Detecting Network Reconnaissance with the Cisco Cyber Threat Defense Solution 1.0

April 9, 2012

Introduction

One of the earliest indicators of an impending network attack is the presence of network reconnaissance. A critical part of a strong security system is the ability to detect reconnaissance early in the threat lifecycle. When reconnaissance is coming from devices on the Internet outside the corporate network, detection and prevention is a fairly straightforward process involving an IPS and firewall. However, when reconnaissance is coming from an infected device that is already inside the enterprise perimeter (where there is far less visibility to such traffic), it can be much more difficult to detect.

To detect reconnaissance inside the network perimeter, the first thing that is needed is clear and complete visibility to network traffic, especially at the user edge, which is the best place on the network to detect internal reconnaissance. When a detection system has visibility to the user edge, it can report on information not normally available further into the infrastructure. This information could include the MAC address of the infected host, and the specific switch and port the device used to connect.

Furthermore, it is important to have visibility to *all* traffic. Some security products can only scale by using sampled NetFlow data. A sampling strategy using average default sampling parameters typically excludes 95% or more of the traffic, seriously reducing product effectiveness. As a general rule, the closer a detection system is to the source of the reconnaissance, the more likely it will have visibility to the traffic. If the NetFlow-enabled system is deployed pervasively throughout the network, including at the user access edge, it will have much better visibility to the network.

The Cisco® Cyber Threat Defense Solution effectively addresses the problem of reconnaissance detection on the internal enterprise network. The solution elegantly solves all of these difficult deployment challenges by using unsampled hardware-enabled NetFlow pervasively across the network to detect reconnaissance.

Prerequisites

This document assumes the reader has read the Cisco Cyber Threat Defense Solution 1.0 Overview, Design and Implementation Guide, and the Introduction to Cisco Cyber Threat Defense ”how-to” document. Readers will gain the maximum benefit from the examples in this guide if they have installed a fully functioning Cyber Threat Defense Solution, including a switch and router infrastructure that is properly configured for sending NetFlow, a fully functioning Cisco Identity Services Engine environment, and a StealthWatch® Flow Collector and StealthWatch® Management Console. With these in place, security practitioners should then plan on following the step-by-step examples while in front of the StealthWatch® console.

Solution Components

The Cisco Cyber Threat Defense Solution 1.0 is composed of three integrated components:

NetFlow data generation devices. NetFlow is the de facto standard for acquiring IP operational data. Traditional IP NetFlow defines a flow as a unidirectional sequence of packets that arrive at a router on the same interface or sub-interface and have the same source IP address, destination IP address, Layer 3 or 4 protocol, TCP or UDP source port number, TCP or UDP destination port number, and type of service (ToS) byte in their TCP, UDP, and IP headers, respectively.

Flexible NetFlow is the next generation in flow technology and is a particularly valuable component of the Cisco Cyber Threat Defense Solution 1.0. Flexible NetFlow optimizes the network infrastructure, reducing operation costs and improving capacity planning and security incident detection with increased flexibility and scalability.

NetFlow can be enabled on most Cisco switches and routers, as well as some Cisco VPN and firewall devices. In addition, select devices now employ special hardware acceleration, ensuring that the NetFlow data collection process does not impact device performance. This enables NetFlow data collection pervasively throughout the network—even down to the user edge—so that every packet from every network segment and every device is completely visible.

Cisco Identity Services Engine. The Cisco Identity Services Engine delivers all the necessary identity services required by enterprise networks—AAA, profiling, posture, and guest management—in a single platform. In the context of the Cisco Cyber Threat Defense Solution 1.0, the Identity Services Engine can be deployed as either a network appliance or virtual machine and answers the "who" (user), "what" (device), and "where" (which NetFlow-enabled device) questions that tie network flow data to the actual physical network infrastructure.

In an enterprise deployment, the Identity Services Engine provides the central policy enforcement needed to govern a network. The Identity Services Engine can provision and deliver cross-domain application and network services securely and reliably in enterprise wired, wireless, and VPN environments. This policy-based service enablement platform helps ensure corporate and regulatory compliance, enhances infrastructure security, and simplifies enterprise service operations. The Identity Services Engine can gather real-time contextual information from the network, users, and devices and make proactive governance decisions by enforcing policy across the network infrastructure.

Lancope® StealthWatch® system. This NetFlow visibility, network performance, and threat detection solution provides an easy-to-use interface that enables both monitoring and detailed forensics. The solution is composed of two core components: the StealthWatch® Management Console and one or more StealthWatch® FlowCollectors. Additional optional components include a StealthWatch® FlowSensor and a StealthWatch® Flow Replicator.

Network Reconnaissance: A Deeper Look

Unless an attacker already has intimate knowledge of a network, one of the first tasks they will want to perform is to discover the devices on the network to determine which ones can be exploited or attacked. There are many different ways to accomplish this.

One of the earliest reconnaissance methods was simply to sequentially "ping" every IP address on a network, starting with the local subnet, and then expand outward. If an IP address responded to a ping, the attacker knew there was a device active at that IP address, and would add it to a locally list of potential attack targets. This "ping" method would require the attacker to "guess" what subnets existed on the network.

Note: In this document, the terms “endpoint,” “computer,” “host,” and “device” are used interchangeably to mean any network-attached system.

This proved to be a very noisy method of identifying vulnerable hosts. Every time a ping was sent to a subnet, the router for that subnet would generate a Layer 2 Address Resolution Protocol (ARP) request for the target IP address. A router-generated ARP request says: “I have someone looking for IP address a.b.c.d. If you are that IP address, please respond with your MAC address so I can forward this packet to you.” A host at that IP address would send an ARP reply that included its MAC address, and then the router would use that MAC address to forward the packet(s). However, if a ping "missed" a target IP address because it was not active on a network, no ARP reply would occur.

On most networks, a sizable amount of the IP address space is unused. A device performing reconnaissance would generate a large number of ARP requests but receive fewer ARP responses, which would result in an ARP “imbalance.” Simple tools that looked for ARP requests with no ARP replies could be used to detect network reconnaissance. Over time, this simple capability of looking for unfulfilled ARP requests has proven to be a reliable method of reconnaissance detection.

To adapt, attackers have become more sophisticated. They have learned to slow down the speed of their reconnaissance in order to "hide" the reconnaissance in

background “noise” of the local network, to make it indistinguishable from other network activity. This significantly slows down the process of network reconnaissance, but it improves an attacker’s chance of not being detected.

Another strategy is to send other kinds of packets, such as UDP or TCP packets, to highly randomized addresses. Instead of sequentially sending packets to each IP address in a range, hackers randomize the destinations and disguise them to look more like normal network traffic. Using specialized software, an attacker can even bypass the normal network stack to send custom crafted packets—packets that might, for example, contain illegal flag conditions such as "SYN/FIN."

This technique provided a way to evade IDS products that had not adapted to this possibility, and allowed the attacker to study the responses. Attackers learned that different network stack vendors would respond to these strange flag conditions in different ways. If you had a list of how each operating system would respond, an attacker could "guess" with fairly high accuracy the kind of host at the other end—Windows, Unix, Linux, Mac OS, etc.—and would allow the infected host to target only those types of devices most vulnerable to attack.

In summary, network reconnaissance can take many different forms. Being able to recognize these forms is key for early detection of network-based threats.

Scenario Overview

The Cisco Cyber Threat Defense Solution 1.0 addresses the problem of network reconnaissance by detecting common reconnaissance patterns in NetFlow data and alerting security administrators when those patterns occur. Cisco's solution integrates StealthWatch® with Cisco's hardware-supported NetFlow and the Identity Services Engine to provide a convenient and effective way to detect network reconnaissance.

In this document, we explore a use case for detecting network reconnaissance. This scenario involves searching for a rogue device emitting events that StealthWatch® categorizes as *scans* (reconnaissance). This use case shows the reader how to locate scan events using StealthWatch® and how to determine both the host responsible for generating the events and the user identity logged into the infected device.

A Note About Concern Indexes

The Cisco Cyber Threat Defense Solution 1.0 provides the ability to detect reconnaissance based on the analysis of NetFlow data. The StealthWatch® console that is used to view the NetFlow data uses a technology called a *concern index* to reflect the severity of a security event. A *concern index* is a numeric value—a counter of sorts—indicating how many times a specific kind of event has occurred within a window of time. The StealthWatch® detection engine examines each flow as it enters the FlowCollector and then applies a set of rules to each flow. The result of the comparison between the rule and the flow data determines whether various counters should be increased.

Independent of the collection process, StealthWatch® constantly analyzes these flows to determine if thresholds have been exceeded or suspicious patterns have been detected. When a concern index value exceeds a defined threshold, StealthWatch® raises an alarm, indicating a potential problem. In this document, we explore using concern indexes to detect network reconnaissance.

Procedure 1: Use StealthWatch® to Find Network Reconnaissance

In this procedure, we use StealthWatch® to identify hosts that are generating concern index alarms. This is the first step in identifying reconnaissance activity.

Step 1 From the main console screen, select the host group *Inside Hosts.*

Step 2 Right-click the host group.

Step 3 From the resulting popup menu, select Security ➔ Concern Index.

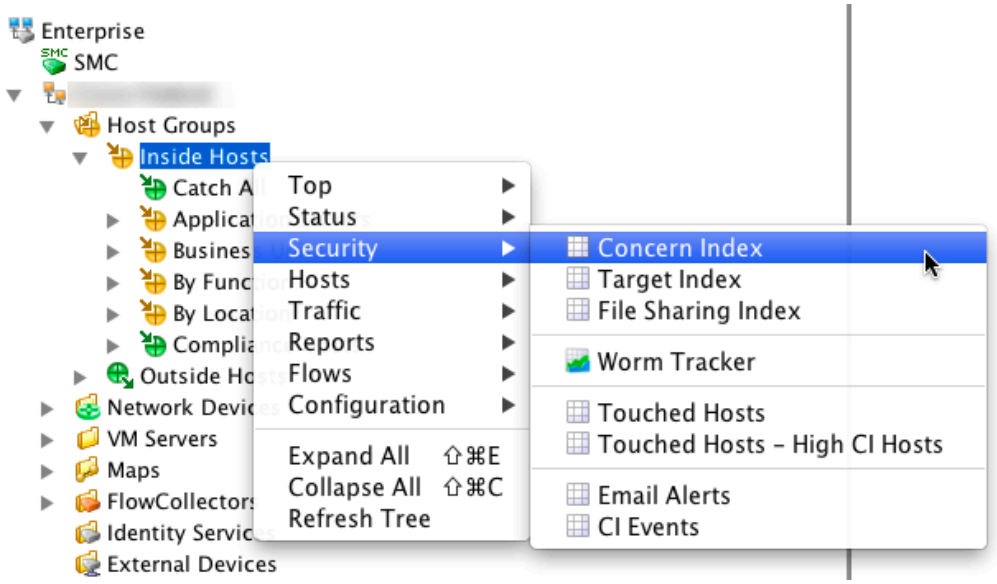

The screen below shows select columns of the resulting concern index table indicating inside hosts that are generating concern index alarms and the underlying cause of the alarms.

Tip: The value in the *CI* column is the concern index. You'll notice the number may range from the low tens or hundreds of thousands to high hundreds of millions. The concern index should be viewed in the context of the *CI%* (*concern index percentage)*. The concern index percentage is the *percentage* of events above or below the *concern index threshold*. These are sorted in descending order on the table, and are color-coded for severity. A high number does not necessarily mean it is the worst problem—the *CI%* should be used to determine severity, not a raw concern index score.

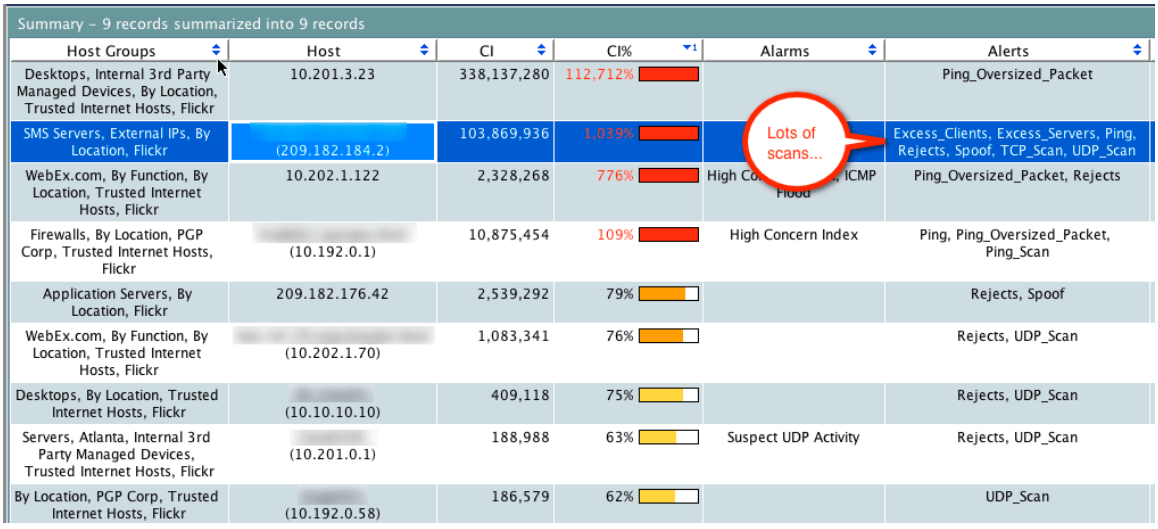

Procedure 2: Examine Details of Network Reconnaissance

In this procedure, we take a deeper look at the details of the network reconnaissance, inspecting the targets of the reconnaissance and the types of communication involved.

Step 1 From the concern index table above, select a host that has scans listed as a type of alert in the alert column. In this example, we select the second record because it has a wide range of alerts, including pings, TCP scans, and UDP scans.

Tip: Each element (cell) in the table is *context-sensitive*. You must be specific about which element in the row you click on, because right-clicking on different elements in the row produces different results.

Step 2 Right-click the host IP address.

Step 3 From the resulting popup menu, select *Host Snapshot*.

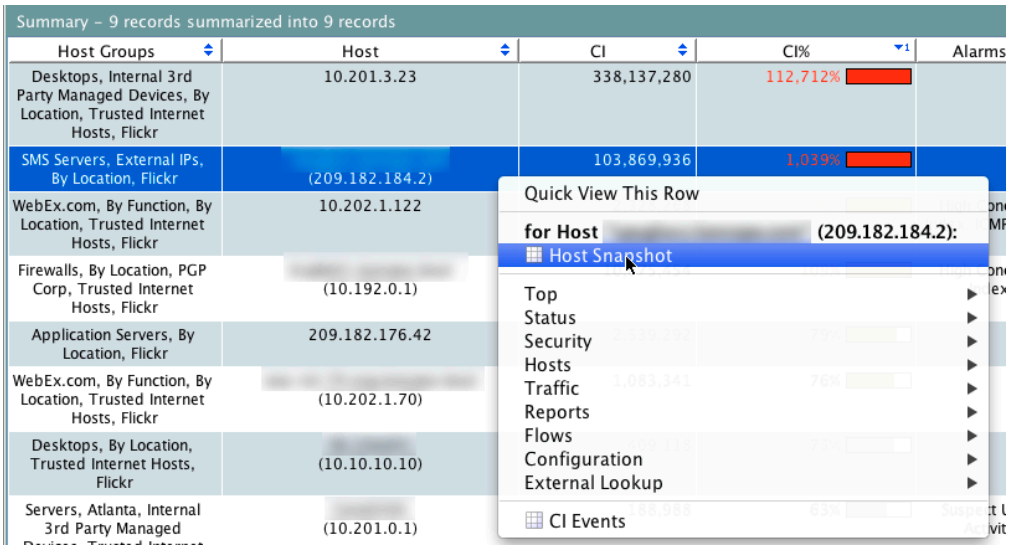

Step 4 On the resulting screen, select the *Top Active Flows* tab.

The resulting table has some interesting characteristics. First, you will notice that there is a high variance of IP addresses with which this source host is communicating. Next, you will notice some of the destinations resolve to places like the Russian Federation. Finally, there are many different ports (services) used in the communication, including some that are unidentifiable. All of these facts were used by StealthWatch® to compute this host's concern index and determine it was scanning other hosts.

Procedure 3: Graph the Reconnaissance Activity

Now that we have a clearly identified host responsible for network reconnaissance, we graph the host activity to look for large-scale patterns to better understand its behavior over time.

Step 1 From the concern index table, right-click on the host and select Flows ➔ Peer vs. Port.

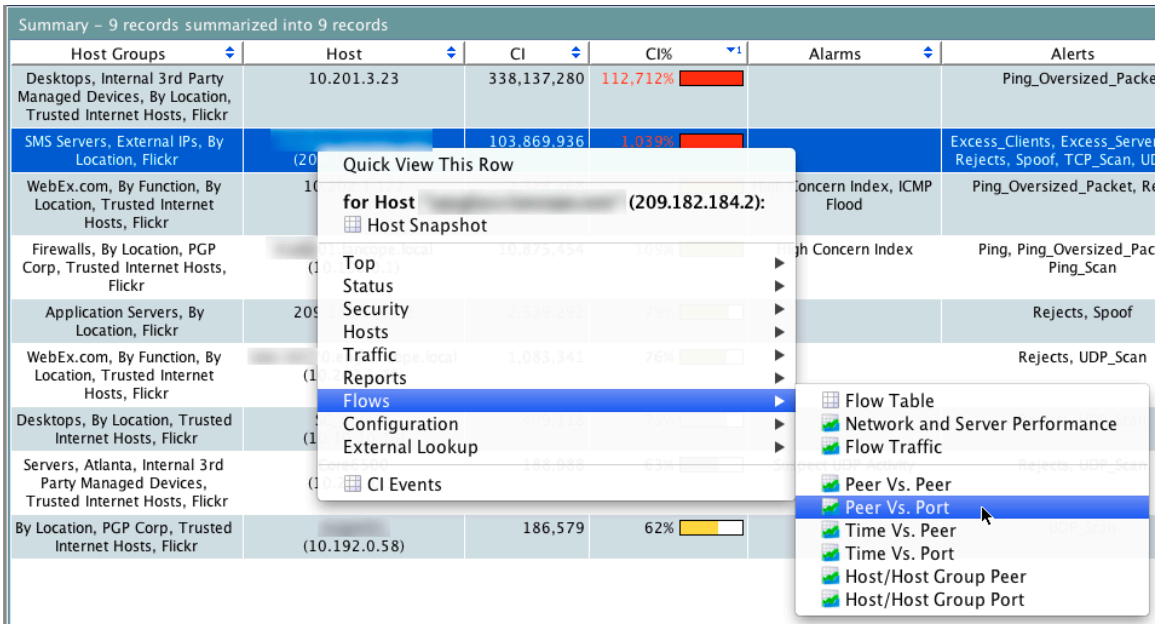

This produces the following graph.

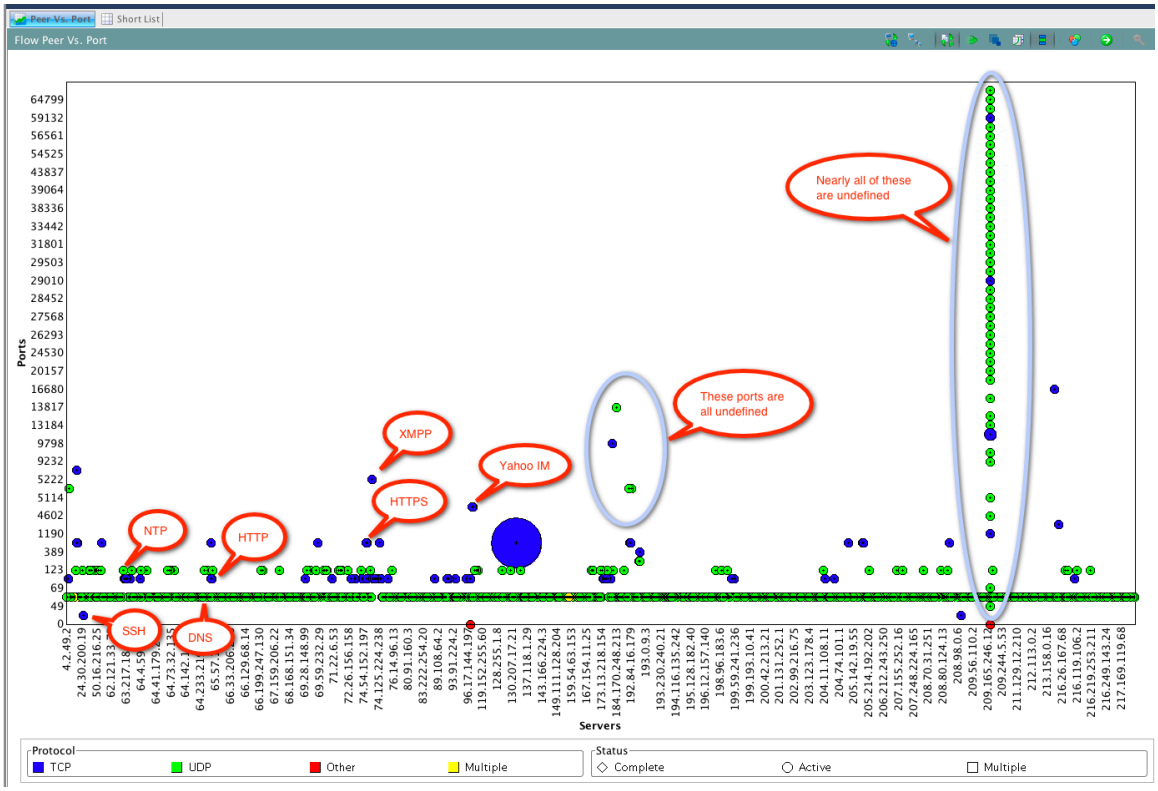

This *Peer vs. Port* graph warrants a close look. It maps all flows during a specified time range against the ports they used to communicate. On the X-axis are IP addresses that are acting as servers. On the Y-axis are the service port numbers used in the communication. The long horizontal green line of dots is made up of DNS requests made by hosts for the IP address associated with the target IP address (the server). For example, when a request is made to https://www.google.com, it results in a DNS request and is reflected as a green dot in this graph. You'll also see SSH, NTP, HTTP, HTTPS, XMPP, and IM all over this chart.

Notice the long vertical line of green dots? It looks out of place, doesn't it? This is because that line is caused by the IP address that generated the scan alert we're investigating: 209.182.184.2. The fact that most of the dots are green and not blue means they are UDP scans, and the fact that there are so many means the source host is generating (probably) random ports for each flow. In other words, it is just looking for hosts with random open ports. This is network reconnaissance shown in an easy to read graph that leaves nothing to the imagination.

Tip: Whenever you use graphs to inspect flows, you are typically looking for patterns that would not be obvious in an unsorted table of values. Using a graph is the easiest way to spot this sort of pattern: Typically, it will look like a long horizontal or vertical line, and will be easy to spot once you know what to look for.

Procedure 4: Identify the User ID of the Infected Host

At this point, we've identified a host that is clearly performing network reconnaissance. Normally, the next step would require trying to hunt down the device with the IP address we've identified. However, because the Cisco Cyber Threat Defense Solution 1.0 integrates with the Cisco Identity Services Engine, the name of the logged-in user ID is also reported in StealthWatch®, making identification of the host a simple process.

Step 1 From the concern index table, right-click on the host and select *Host Snapshot* from the popup menu.

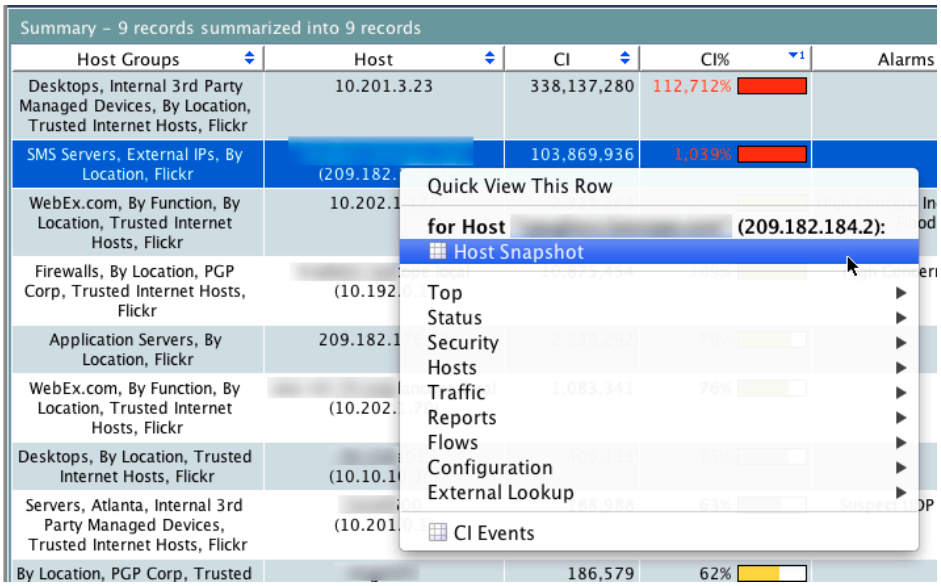

This produces an identity and device table with the *Identity, DHCP & Host Notes* tab selected. A portion of that table is listed below showing the user ID of the individual logged into the infected device. You will also notice the Identity Services Engine reports a plethora of additional information about the host, including the MAC address of the device (including the manufacturer), the identity group the device belongs to, and the device type.

Cisco ISE	User Name	MAC Address	Identity Group	VLAN	Device Type
(10.35.48.243)		00:22:68:1a:59:d0 (Hon Hai Precision Ind. Co., Ltd.)	Profiled:Workstation		Windows7-Workstation

At this point, it is a simple matter to track down the user and their device for remediation.

Conclusion

In this document, we used the Cisco Cyber Threat Defense Solution 1.0 to identify a host with a particularly high concern index. We discovered the host was exhibiting unusual behavior, then looked more closely at that behavior and determined that the host was performing reconnaissance on the network and was probably infected. Using the Cisco Cyber Threat Defense Solution 1.0, we were able to quickly identify this malicious host.