Global Information Assurance Certification Paper

I Can See You Behind Layer 2…

Overcoming the difficulties of Packet Capturing on a Switched Network

A SANS/GIAC Security Essentials Certification Practical Assignment

May 21, 2003

By:

Douglas C. Hewes

OVERVIEW

In the past most networks utilized concentrators or hubs to connect various clients, servers and other network hosts. These are known as Layer-1 devices, which operate solely on the Physical Layer of the OSI model. These devices worked well for small networks. Essentially a packet went from a networked computer to the hub and then out to all other networked devices. This created a single collision domain, where all devices were able to see all traffic; regardless of whether the traffic was intended for that particular host or not. When in normal operating mode, the network adapters ignored all packets not addressed to them. It wasn't long before packet-capturing tools were created to take advantage of this design, along with network adapters that supported promiscuous mode. These utilities were, and still are, extremely useful in detecting and resolving network communications issues. Quite frequently network administrators must "take to the wire" to figure out exactly what is going on. An unfortunate side effect of this is that unsavory users can utilize these same tools to collect a world of information ranging from clear-text passwords to SMTP traffic to basically anything that goes across the network.

As networking grew more sophisticated, devices were engineered to recognize and utilize layer-2 information, in which each device's Media Access Control (MAC) address resides. The driving reason for this was speed. Since Ethernet is a collision detection, rather then collision avoidance topology, a single collision domain with a large number of nodes led to excessive packet collisions. This inherently limited the effective size of a layer-1 based network. In a network comprised of layer 2 devices, or a "switched network", each layer-2 device builds an Address Resolution Protocol (ARP) table. This table records MAC addresses and which ports those addresses are patched in to. When a packet is transmitted from a node, the switch then compares the destination MAC address against it's ARP table and forwards the packet to only that port. This in effect creates a separate collision domain for each node. Packet collisions are reduces to almost nothing with only broadcasts and packets destined for MAC addresses not in the ARP table being forwarded to all nodes[1]. Although a layer-2 based network is considerably faster and more efficient then a layer-1 network, there is an interesting side effect. Packet-capturing sniffers now could only see traffic to and from the sniffer itself. This eliminated the administration benefits of sniffers, but also greatly increased security since rogue sniffers were no longer effective at capturing conversations and passwords not meant for them.

To overcome this limitation there are several high-end tools, such as Network Observer, that claim to be able to capture traffic from all ports. Most switches also have some level of manageability whereby you can use "port-mirroring" to copy all traffic destined for say the server port to a sniffer port. In most legitimate installations of sniffer utilities, remote agents are patched down into each switch in the network, port mirroring is utilized and the captures are sent to a main console. Then there are the less the legitimate installations. This paper will provide a high level overview of the tools and techniques that a hacker or unsavory internal user, may use to capture traffic on your layer-2 network.

OSI LAYER 1 AND 2 CONCEPTS

The Open Systems Interconnection Reference Model, or OSI Model, categorizes network communications, protocols and applications into 7 layers. The lowest layer, layer-1, is the Physical layer. This layer is essential concerned with the physical, mechanical and electrical connectivity between network cards in nodes. There is no possessing or decision making at this level – just the transmission of ones and zeros[2]. A concentrator or hub at it's lowest level is simply a patch block where all network connections are tied together. In the below diagram, the target client is sending a password (or any data) to the network file server. Note that in this scenario the packet goes to all networked workstations.

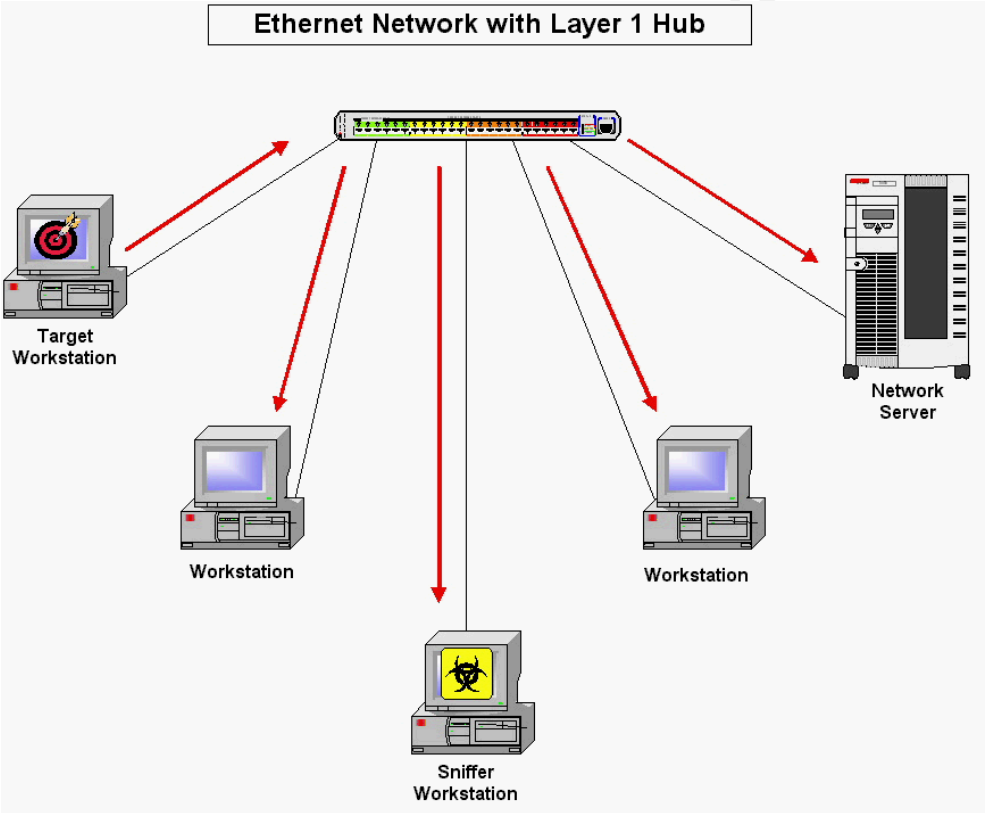

The next layer in the OSI Model, Layer-2, is the Data Link Layer. At this layer the MAC addresses for each node is taken into account. Other functions, such as error control and bit stuffing also occur at this layer[3]. Layer-2 devices such as bridges and switches use these MAC addresses to forward traffic only to the intended recipient. Over time switches build an internal network topology map in the form of ARP tables. This table records all of the MAC addresses for devices patched into the switch. This table is refreshed every 5 minutes or so, but that time is generally customizable. Through the use of this ARP table, packets are sent only from the source to the destination. All other nodes are unaware of any communications. The only exceptions to this are broadcast packets, and packets that are sent to nodes not listed in the local ARP table. These packets are generally forwarded to all nodes. This behavior is somewhat

customizable on some switches however. Since all nodes have a virtual "clean" line to and from the switch, packet collisions are almost completely eliminated. Note the same packet being transmitted from the target workstation to the network server is no longer sent to every other workstation on the network:

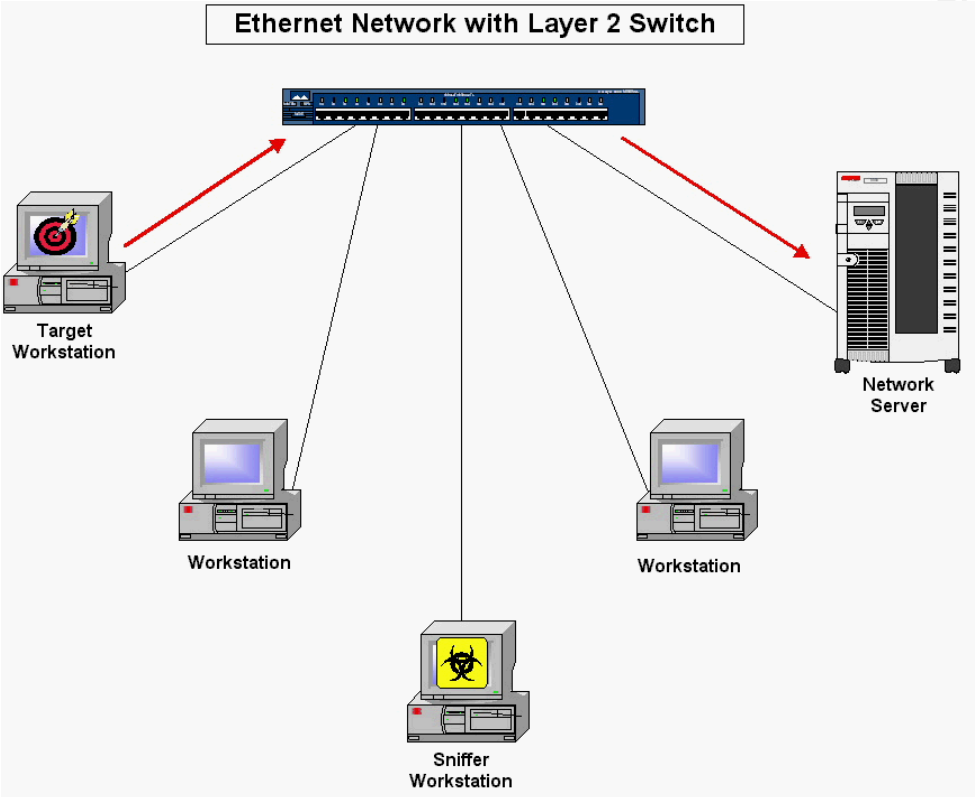

The sniffer workstation, either legitimate or subversive, can no longer intercept and record the traffic.

ARP SPOOFING-INTERNAL BRIDGING TECHNIQUES

During the early stages of an attack, a hacker (either inside or outside the network) must perform some level of recon or network discovery. During this phase such useful information as usernames, workstation names, servers, protocols, addresses and even passwords are collected for later use. The primary method of accomplishing this is the packet-capturing sniffer. The switched network prevents this…normally. In order the overcome this obstacle, the preferred method is ARP Spoofing. Recall that the switch relies on the information it previously collected to know where to forward packets. This table not only refreshes periodically, but on some switches also welcomes un-requested updates. This is when the ARP spoofing tools craft an ARP reply packet that provides the switch with false information. Some tools, such as van Hauser / THC's parasite v.0.5 (http://www.infowar.co.uk/thc/) have been adapted to wait until the switch sends out an ARP request, and thereby bring less suspicion on themselves. The goal with this attack is to trick a switch into thinking that a server or workstation has moved to another port. The switch then

forwards all traffic destined for the legitimate target to the sniffer workstation. The obvious limitation here is that the person using the ARP spoofer has to have discovered the real MAC address of the target, and can only capture a conversation with a single node at a time. If that node happens to be the file servers though, that is all an attacker needs. The second portion of this attack is to then maintain the originally intended conversation by acting as a bridge and forwarding the packet on to the real destination[4]. In this manner the connection between the original host and destination is not broken. Some tools, such as ARP0c2 will do this automatically. Other tools will utilize fragrouter for this purpose[5]. This is commonly known as the "man-in-the-middle" attack.

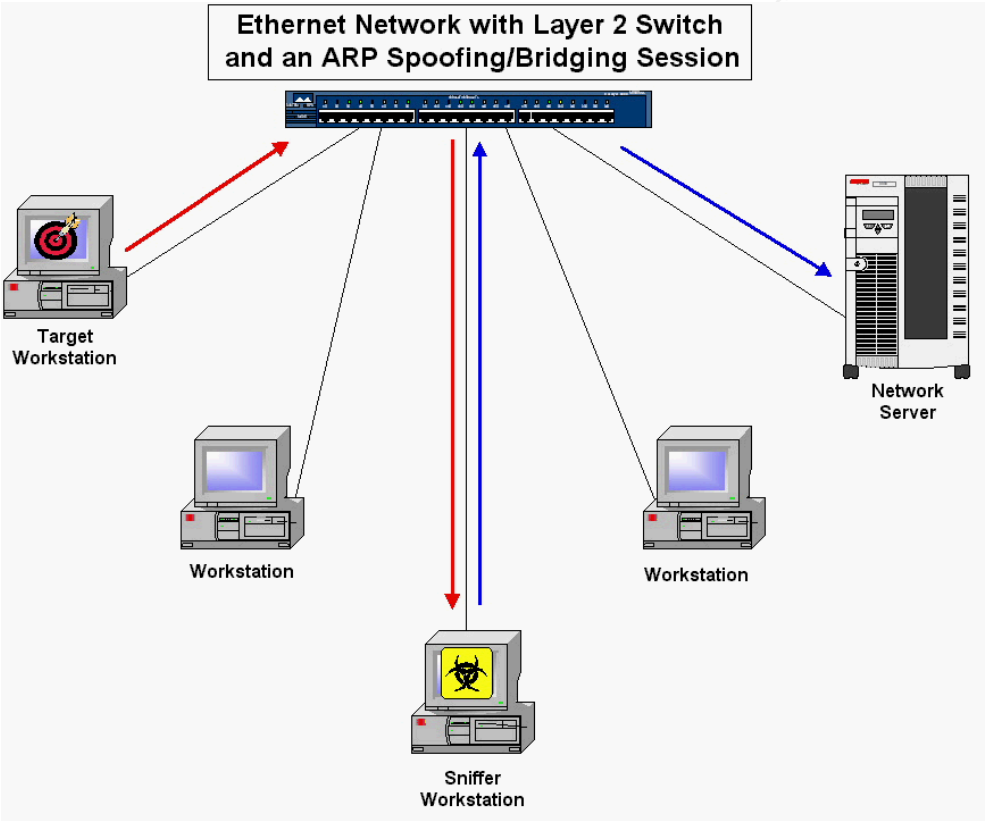

MAC ADDRESS OVERFLOW TECHNIQUES

Switches have a finite amount of memory. This varies from switch to switch, but essentially all switches have a limit to the maximum number of MAC addresses they can record in their ARP tables. Under normal circumstances this number is many times greater then the number of ports on the switch, and therefore is seldom a concern. Someone trying to overcome the Layer-2 issue may utilize this limitation to his advantage. By generated a large number of MAC addresses and force-feeding them to the switch, a buffer overflow type condition will be created. During this period the switch cannot keep up with maintaining an ARP table and forwarding packets correctly. Most switches will then fail back to a Layer-1 mode. In essence, the switch will begin acting like a hub; forwarding

all traffic to all ports. The limitations with this technique are that the switch and the network itself will be loaded to a much higher level then normal and performance will suffer. Therefore this technique is seldom used for long periods of time. It is more often used long enough to gather the MAC addresses needed to perform the aforementioned ARP Spoofing technique. TCH-Parasite also includes this type of functionality[6].

OTHER MEANS

Although these are the most common methods to circumvent the additional security offered by Layer-2 devices, there are others.

Physical Access

To begin with, someone with physical access to a switch can almost always enable port mirroring. With physical access someone can attach a management cable, reset a switch to erase passwords, establish port mirroring with any or all ports and even set up their own password. At this point they can telnet in and mirror any port at any time.

Switch Passwords

The author still finds that most networks may secure routers and servers and even workstations with passwords, but more often then not switches are not secured. By default most switches can be managed through telnet sessions – with more and more supporting internal HTTP stacks allowing for easy web browser based management. Without setting a password and ensuring physical security, the network is vulnerable.

MAC Address Duplication

On a Linux workstation (and others) it is possible to manually set the reported MAC address of your Ethernet card. This is done with the following commands[6]:

```
ifconfig eth0 down
ifconfig eth0 hw ether 01:01:01:01:01:01
ifconfig eth0 up
```

This will not work on some switches that will sense the duplication – something that should never occur under normal circumstances. This fact will not stop someone from trying it though.

Make Them Come to You

If your attacker's goal is password sniffing there are methods to force users to try to authenticate to his workstation. One such simple method is to send out an email to all users in an NT domain asking them to download the

latest virus information (or anything enticing enough to get them to click the link). Then provide a link to a file on a shared drive in this format:

file:////yourcomputer/sharename/message.html

If the attacker is running a tool such as L0pht Crack with SMB capture on then he will effectively capture your username and password hash for later cracking. Although this isn't a way to directly circumvent Layer-2 technology, it is an effective means in a switched network environment[7].

AVAILABLE TOOLS & INTERNET RESOURCES

ARP0c2 (Linux) and WCI (Windows)
ARP Spoofing tool
http://www.phenoelit.de/arpoc/index.html

dsniff (Linux)
A suite of tools including ARP Spoofing, MAC flooding and more
http://naughty.monkey.org/~dugsong//dsniff/

THC-Parasite v.0.5 (Linux)
Includes both ARP Spoofing and MAC Flooding
http://www.infowar.co.uk/thc/ or
http://packetstorm.securify.com/sniffers/parasite-0.5.tar.gz

smit (Linux)
ARP Spoofing and ARP Query utility
http://packetstorm.securify.com/sniffers/smit.tar.gz

Additional tool found that a Packetstorm search:
http://209.143.242.119/cgi-bin/search/search.cgi?searchvalue=switched+sniffer&type=archives

PREVENTATIVE MAINTENANCE

Keep in mind that the Layer-2 functionality of a switch was never intended as a means to securing a network. It was designed to increase the speed and efficiency of the network. There are a few steps that can be taken to reduce the chances of an attacker using these techniques. First secure physical access to all switches. Next make sure that strong passwords are in place on all switches. Remove telnet and HTTP management on switches if not absolutely needed. Some switches, such as Cisco Catalyst switches, support Port Secure. This enables an administrator to lock in a specified MAC address for each port. If you switch supports this, consider using it. Ultimately, however, only the use of encryption can really prevent the threat of a rogue sniffer.

CITED SOURCES

[1] Cisco Systems. "Introduction to LAN Protocols." Internetworking Technology Overview. June 1999. URL: http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/introlan.htm

[2] Briscoe, Neil. "Understanding the OSI 7 Layer Model." July 2000. URL: http://www.itp-journals.com/search/t04124.htm

[3] TechTarget.com. "OSI." Computer Fundamentals: Standards and Organizations. 27 July 2000. URL: http://www.whatis.com/WhatIs_Definition_Page/0,4152,212725,00.html

[4] FX of Phenoelit. "ARP0c/WCI ARP Interceptor." 2000. URL: http://www.phenoelit.de/arpoc/index.html

[5] McClure, Stuart & Scambray, Joel. "Switched networks lose their security advantage due to packet-capturing tool." *InfoWorld*. 26 May 2000. URL: http://www.infoworld.com/articles/op/xml/00/05/29/000529opswatch.xml

[6] van Hauser / THC. "THC-Parasite v.0.5 README." 2000. URL: http://www.infowar.co.uk/thc

[7] L0pht Heavy industries. "L0phtCrack 2.5 FAQ." 12 October 1999. URL: http://www.l0pht.com/l0phtcrack/faq.html