Packet Sniffing on Layer 2 Switched Local Area Networks

Ryan Spangler
ryan@packetwatch.net
Packetwatch Research
http://www.packetwatch.net

December 2003

Abstract

Packet sniffing is a technique of monitoring network traffic. It is effective on both switched and non-switched networks. This paper discusses several methods that result in packet sniffing on Layer 2 switched networks. Each of the sniffing methods will be explained in detail. The purpose of the paper is to show how sniffing can be accomplished on switched networks, and to understand how it can be prevented.

Table of Contents

1.0 Introduction

Packet sniffing is a technique of monitoring network traffic. It is effective on both switched and non-switched networks. In a non-switched network environment packet sniffing is an easy thing to do. This is because network traffic is sent to a hub which broadcasts it to everyone. Switched networks are completely different in the way they operate.

Switches work by sending traffic to the destination host only. This happens because switches have CAM tables. These tables store information like MAC addresses, switch ports, and VLAN information [1]. Before sending traffic from one host to another on the same local area network, the host ARP cache is first checked. The ARP cache is a table that stores both Layer 2 (MAC) addresses and Layer 3 (IP) addresses of hosts on the local network. If the destination host isn't in the ARP cache, the source host sends a broadcast ARP request looking for the host. When the host replies, the traffic can be sent to it. The traffic goes from the source host to the switch, and then directly to the destination host. This description shows that traffic isn't broadcast out to every host, but only to the destination host, therefore it's harder to sniff traffic.

This paper discusses several methods that result in packet sniffing on Layer 2 switched networks. Each of the sniffing methods will be explained in detail. The purpose of the paper is to show how sniffing can be accomplished on switched networks, and to understand how it can be prevented.

2.0 Attacks

This paper shows that several attacks are available to sniff layer 2 switched networks. These attacks that result in sniffing are: ARP cache poisoning, CAM table flooding, and switch port stealing. Each attack will be explained in detail.

2.1 ARP Cache Poisoning

This attack uses Address Resolution Protocol (ARP) spoofing to sniff traffic between hosts. ARP spoofing is possible because of the exploitation of gratuitous ARP. Gratuitous ARP is when an ARP reply is sent without first receiving an ARP request [2]. ARP cache poisoning works by poisoning the ARP cache of the target hosts [3]. The attacker wanting to sniff the traffic essentially inserts his computer between the target hosts and forwards traffic back and forth between computers. The way the attack works will be explained below (see Figure 1).

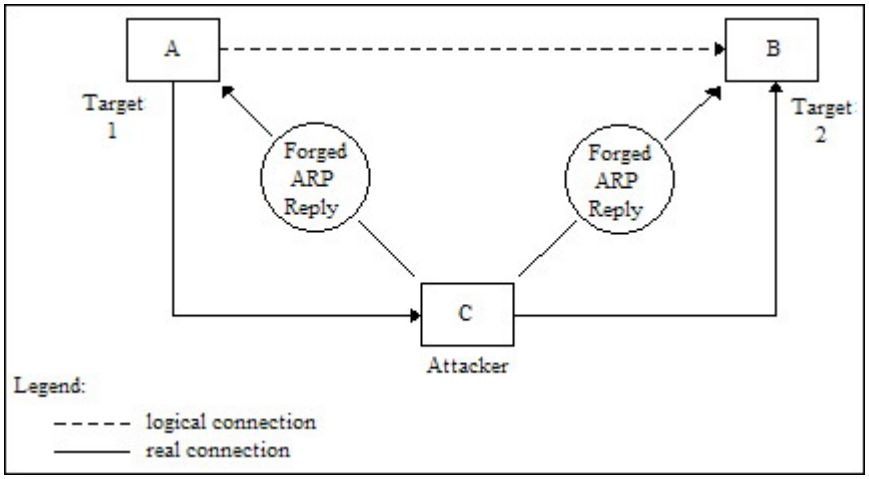

Figure 1: ARP Cache Poisoning

The attack starts by having the attacker send a forged gratuitous ARP packet with host B's IP address and the attackers MAC address to host A. The attacker also sends a forged gratuitous ARP packet with host A's IP address and the attackers MAC address to host B. Now, all of host A and host B's traffic will go to the attacker, where it can be sniffed, instead of directly to each other.

2.2 CAM Table Flooding

This attack uses MAC flooding to sniff traffic on the local area network. Content Addressable Memory (CAM) table flooding works by flooding the CAM table. CAM tables store information like MAC addresses, and switch ports, along with their VLAN information. CAM tables have fixed sizes, so they can only store a certain number of entries. The user wanting to sniff the traffic floods the switch with MAC addresses until the CAM table is full, at which point the switch starts to broadcast the traffic [4].

The attack starts by having the attacker flood the network with forged gratuitous ARP packets that each contains unique source MAC addresses. This causes some switches to go into a hub-like mode forwarding all traffic to all ports. What happens is that once the CAM table is full, the traffic without a CAM entry floods on the local VLAN. The already existing traffic with existing entries in the CAM table will not be forwarded out on all of the ports. Now, with the traffic being broadcasted to everyone, there will be no trouble sniffing it.

2.3 Switch Port Stealing

This attack uses MAC flooding to sniff traffic between two hosts. Switch port stealing works by stealing the switches port of the target host. Switches learn to bind MAC addresses to each port by seeing the source MAC addresses in the packets that arrive from each port. The user wanting to sniff the traffic steals the switches port to the target host so the traffic will go through it first, then to the target host [5]. The way the attack works will be explained below (see Figure 2).

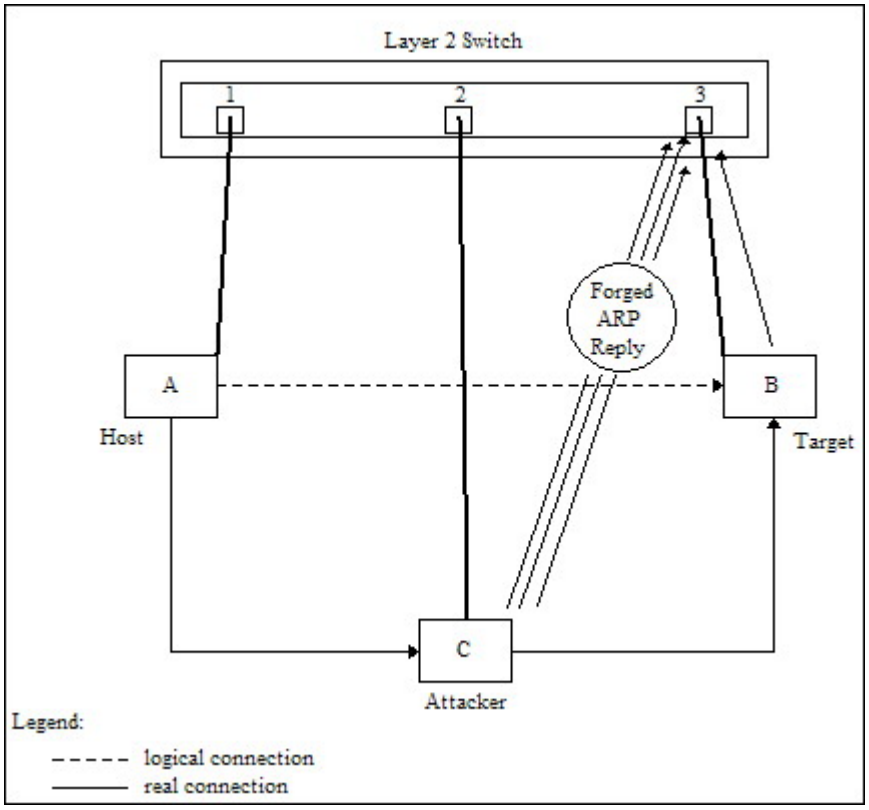

Figure 2: Switch Port Stealing

The attack starts by having the attacker flood the switch with forged gratuitous ARP packets with the source MAC address being that of the target host and the destination MAC address being that of the attacker. The flooding process described here is different than the flooding process used in CAM table flooding. Since the destination MAC address of each flooding packet is the attackers MAC address, the switch will not forward these packets to other ports, meaning they will not be seen by other hosts on the network. Now, a race condition exists because the target host will send packets too. The switch will see packets with the same source MAC address on two different ports and will constantly change the binding of the MAC address to the port. Remember that the switch binds a MAC address to a single port. If the attacker is fast enough, packets intended for the target host will be sent to the attacker's switch port and not the target host. The attacker has now stolen the target hosts' switch port. When a packet arrives to the attacker, the attacker performs an ARP request asking for the target hosts' IP address. Next, the attacker stops the flooding and waits for the ARP reply. When the attacker receives the reply, it means that the target hosts' switch port has been restored to its original binding. Now, the attacker can sniff the packet, then forward it to the target host and restart the flooding process waiting for new packets.

3.0 Defenses

There are several ways to mitigate these packet sniffing attacks. The first of these actions is to enable port security on the switch. Port security is a feature found on high-end switches that ties a physical port to a MAC address. This allows you to either specify one or more MAC addresses for each port, or learn a certain number of MAC addresses per port. A change in the specified MAC address for a port or flooding of a port can be controlled in many different ways through switch administration.

An important fact to know is that port security capabilities are dependant on the platform meaning that different switch manufacturers have different capabilities.

The second way to mitigate sniffing is through the use of static ARP entries. Static ARP entries are permanent entries that won't time out from the ARP cache [2]. This method does have a drawback though. Administrators have to create new entries on every host on the network every time a new host is connected, or when a network card is replaced.

The final method of defense is through detection. Intrusion detection systems can be configured to listen for high amounts of ARP traffic. There are also tools specifically designed to listen for ARP replies on networks [6]. This method is prone to reporting false positives though. It should be remembered that detection is always an important step in mitigation.

4.0 Summary

Switched networks are designed to allow computers to communicate more effectively with other local computers. Switches were never intended to be used as security features, as they are today. They are designed to increase the efficiency of available bandwidth on networks. This is accomplished by sending traffic to only the intended destination. Because of this sniffing is harder to do on switched networks. Vulnerabilities found in modern networking protocols have lead to several methods of sniffing though. Fortunately, there are several methods of prevention.

5.0 References

[1] S. Convery, "Hacking Layer 2: Fun with Ethernet Switches", Blackhat [Online Document], 2002, Available HTTP: http://www.blackhat.com/presentations/bh-usa-02/bh-us-02-convery-switches.pdf

[2] W. R. Stevens, TCP/IP Illustrated, Volume 1: The Protocols. Addison-Wesley, 1994.

[3] http://ettercap.sourceforge.net/

[4] http://www.monkey.org/~dugsong/dsniff/

[5] A. Ornaghi, M. Valleri, "Man in the middle attacks Demos" Blackhat [Online Document], 2003, Available HTTP: http://www.blackhat.com/presentations/bh-usa-03/bh-usa-03-ornaghi-valleri.pdf

[6] S. Whalen, "An Introduction to ARP Spoofing", Node99 [Online Document], April 2001, Available HTTP: http://www.node99.org/projects/arpspoof/