Application Note

SAFE Enterprise Layer 2 Addendum

Introduction

The SAFE Enterprise white paper published by Cisco Systems discusses various network attacks on a large-scale enterprise network. These network attacks are based on the premise that each device on the network is a potential target. During the time since the original publication date of the SAFE Enterprise white paper, significant research on network attacks has been conducted, focusing on Layer 2 of the OSI reference model. This research has prompted the need for an update to the white paper focusing on more specific requirements to protect Layer 2 in the network infrastructure.

General Switch Operation

Unlike hubs, switches are able to regulate the flow of data between their ports by creating almost "instant" networks that contain only the two end devices communicating with each other at that moment in time. Data frames are sent by end systems and their source and destination addresses are not changed throughout the switched domain. Switches maintain Content-Addressable Memory (CAM) lookup tables to track the source addresses located on the switch ports. These lookup tables are populated by an address-learning process on the switch. If the destination address of a frame is not known or if the frame received by the switch is destined for a broadcast or multicast address, the switch forwards the frame out all ports. With their ability to isolate traffic and create the "instant" networks, switches can be used to divide a physical network into multiple logical, or virtual LANs (VLANs) through the use of Layer 2 traffic segmentation.

VLANs

VLANs allow network administrators to divide their physical networks into a set of smaller logical networks. Like their physical counterparts, each VLAN consists of a single broadcast domain that is isolated from other VLANs. VLANs work by tagging packets with an identification header and then restricting the ports that the tagged packets can be received on to those that are part of the VLAN. The two most prevalent VLAN tagging techniques are the IEEE 802.1q tag and the Cisco Inter-Switch Link (ISL) tag.

The ISL header format is shown in Figure 1 and the 802.1q VLAN header format is shown in Figure 2. The VLAN header is inserted at Layer 2 and the information contained within the tags is used to identify which VLAN the traffic belongs to. Only those ports that belong to the VLAN specified in the header are capable of receiving the traffic. The destination address then further specifies the particular port within the VLAN where the traffic is destined.

Figure 1
ISL Tag Header Structure

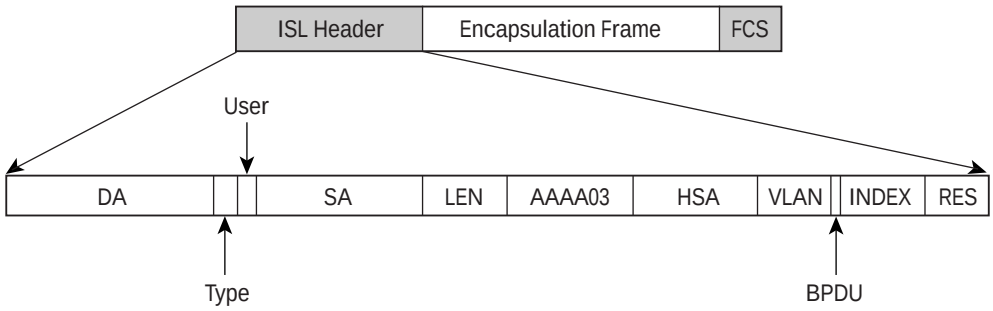

Figure 2
802.1q Header Structure

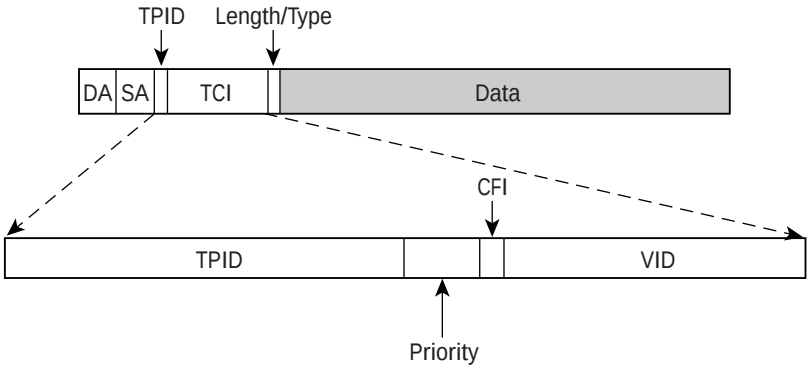

Switches are Targets

Like routers, both Layer 2 and Layer 3 switches have their own sets of network security requirements. But not as much public information is available about the network security risks in switches and what can be done to mitigate those risks. Switches are susceptible to many of the same Layer 3 attacks as routers. . Most of the network-security techniques detailed in the section of the SAFE Enterprise white paper titled “Routers Are Targets” also apply to switches. However, switches, and Layer 2 of the OSI reference model in general, are subject to network attacks in unique ways. These include:

- CAM table overflow
- VLAN hopping
- Spanning-Tree Protocol manipulation
- Media Access Control (MAC) Address spoofing
- Private VLAN
- DHCP “starvation”

This paper explores each of these types of network attacks as well as provides recommendations for how to mitigate or reduce the effects of these attacks. Finally, risks involved with using Cisco Discovery Protocol and VLAN Trunking Protocol are discussed.

CAM Table Overflow

The CAM table in a switch contains information such as the MAC addresses available on a given physical port of a switch, as well as the associated VLAN parameters. When a Layer 2 switch receives a frame, the switch looks in the CAM table for the destination MAC address. If an entry exists for the MAC address in the CAM table, the switch forwards the frame to the port designated in the CAM table for that MAC address. If the MAC address does not exist in the CAM table, the switch forwards the frame out every port on the switch, effectively acting like a hub. If a response is seen, the switch updates the CAM table. Figure 3 shows this operation. In this figure, host A sends traffic to host B. The switch receives the frames and looks up the destination MAC address in its CAM table. If the switch cannot find the destination MAC in the CAM table it then copies the frame and broadcasts it out all of the switch ports (i). Host B receives the frame and sends back a reply to host A. The switch then sees that the MAC address for host B is located on port 2 and writes that information into the CAM table (ii). Now, any frame sent by host A (or any other host) to host B will simply be forwarded to port 2 of the switch and not broadcast out every port as was done earlier (iii).

Figure 3
MAC CAM Table Operation

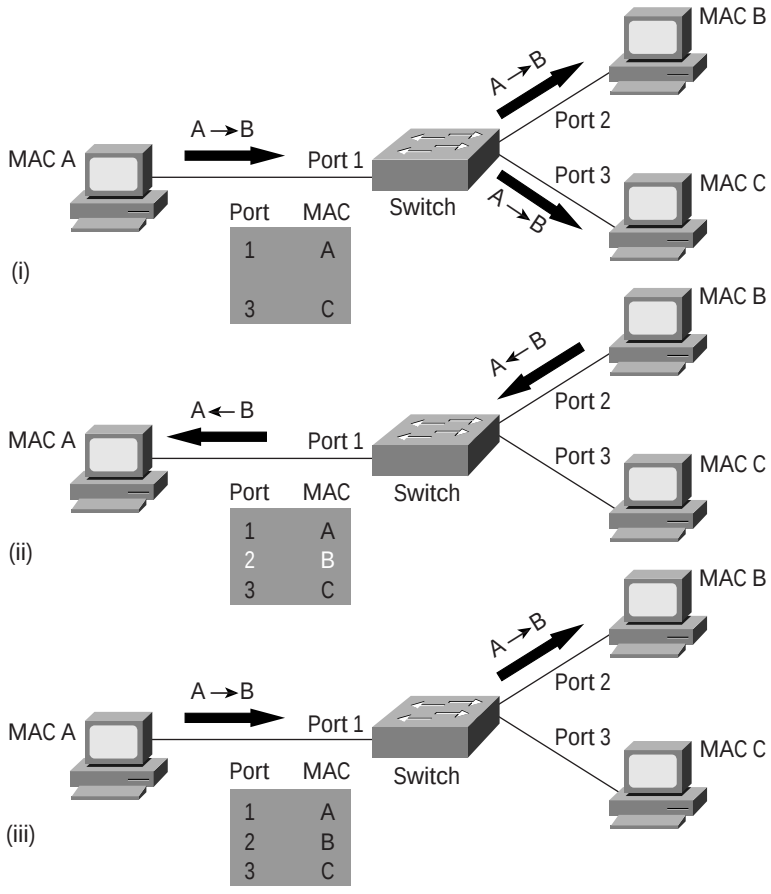

CAM tables are limited in size. If enough entries are entered into the CAM table before other entries are expired, the CAM table fills up to the point that no new entries can be accepted. Typically a network intruder will flood the switch with a large number of invalid-source MAC addresses until the CAM table fills up. When that occurs the switch will flood all ports with incoming traffic because it cannot find the port number for a particular MAC address in the CAM table. The switch, in essence, acts like a hub. If the intruder does not maintain the flood of invalid-source MAC addresses, the switch will eventually time out older MAC address entries from the CAM table and begin to act like a switch again. CAM table overflow only floods traffic within the local VLAN so the intruder will see only traffic within the local VLAN to which he or she is connected.

In May of 1999 the tool *macof* was released. It was written in approximately 100 lines of PERL code and was later ported to C language code and incorporated into the *dsniff* package. This tool floods a switch with randomly generated MAC addresses. When the switch's CAM table fills up with these addresses, the switch begins to forward all frames it receives to every port. Figure 4 illustrates a CAM table-overflow attack. In this figure the host with MAC address C in the bottom right of each frame is sending out multiple packets with various source MAC addresses. Over a short period of time the CAM table in the switch fills up until it cannot accept new entries. When this happens the switch begins to broadcast all packets which it receives out of every port so that packets sent from host A to host B are also broadcast out of port 3 on the switch.

Figure 4
CAM Table-Overflow Attack

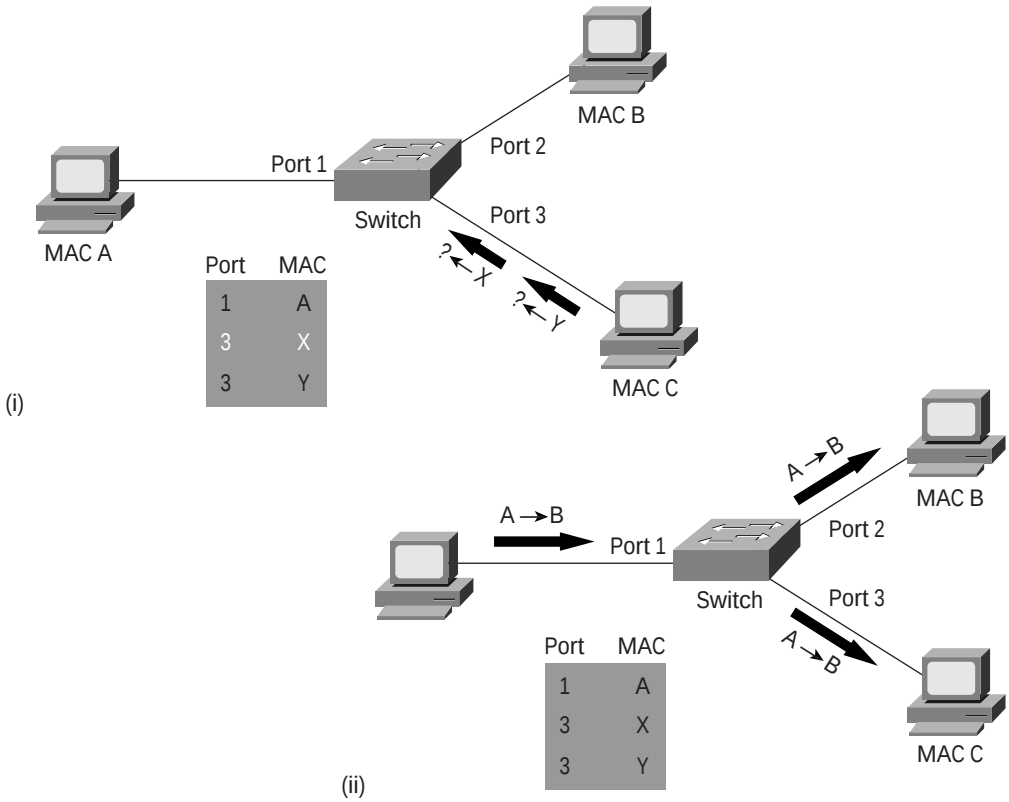

Network Attack Mitigation

The CAM table-overflow attack can be mitigated by configuring port security on the switch. This option provides for either the specification of the MAC addresses on a particular switch port or the specification of the number of MAC addresses that can be learned by a switch port. When an invalid MAC address is detected on the port, the switch can either block the offending MAC address or shut down the port.

Specifying MAC addresses on switch ports is far too unmanageable a solution for a production environment. Limiting the number of MAC addresses on a switch port is manageable.

Command Samples to Mitigate CAM Table-Overflow Attacks
CatOS> (enable) set port security *mod_num/port_num* **enable** *[mac_addr]*
CatOS> (enable) set port security *mod_num/port_range* **enable maximum** *[max_mac_addr]*
CatOS> (enable) set port security *mod_num/port_num mac_addr*
CatOS> (enable) show port *[mod_num[/port_num]*

VLAN Hopping

VLAN hopping is a network attack whereby an end system sends out packets destined for a system on a different VLAN that cannot normally be reached by the end system. This traffic is tagged with a different VLAN ID to which the end system belongs. Or, the attacking system may be trying to behave like a switch and negotiate trunking so that the attacker can send and receive traffic between other VLANs.

Switch Spoofing—In a VLAN hopping attack, the network attacker configures a system to spoof itself as a switch. This requires that the network attacker be capable of emulating either ISL or 802.1q signaling along with Dynamic Trunk Protocol (DTP) signaling. Using this method, a network attacker can make a system appear to be a switch with a trunk port. If successful, the attacking system then becomes a member of all VLANs.

Double Tagging—Another version of this network attack involves tagging the transmitted frames with two 802.1q headers in order to forward the frames to the wrong VLAN (Figure 5). The first switch to encounter the double-tagged frame (1) strips the first tag off the frame and forwards the frame. The result is that the frame is forwarded with the inner 802.1q tag out all the switch ports (2) including trunk ports configured with the native VLAN of the network attacker. The second switch then forwards the packet to the destination based on the VLAN identifier in the second 802.1q header.

Figure 5
VLAN Hopping with Double-Encapsulated 802.1q Traffic

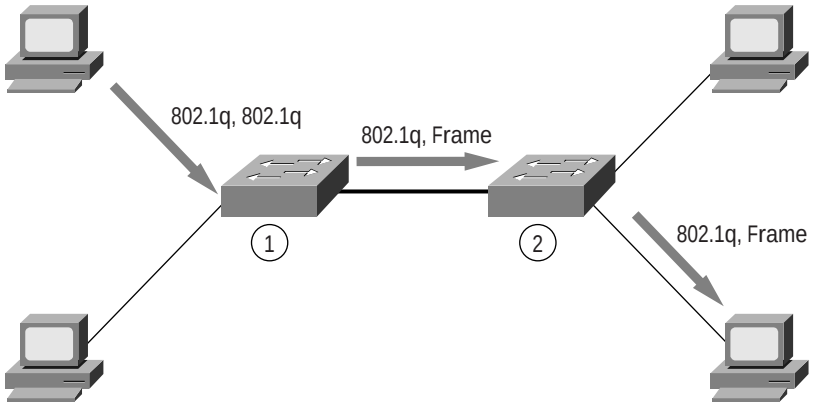

Network Attack Mitigation

Mitigating VLAN hopping attacks requires several modifications to the VLAN configuration. One of the more important elements is to use dedicated VLAN IDs for all trunk ports. Also, disable all unused switch ports and place them in an unused VLAN. Set all user ports to nontrunking mode by explicitly turning off DTP on those ports.

Spanning-Tree Protocol Manipulation

Another attack against switches involves intercepting traffic by attacking the Spanning-Tree Protocol. This protocol is used in switched networks to prevent the creation of bridging loops in an Ethernet network topology. Upon bootup, the switches begin a process of determining a loop-free topology. The switches identify one switch as a root bridge and block all other redundant data paths.

By attacking the Spanning-Tree Protocol, the network attacker hopes to spoof his or her system as the root bridge in the topology. To do this the network attacker broadcasts out Spanning-Tree Protocol Configuration/Topology Change Bridge Protocol Data Units (BPDUs) in an attempt to force spanning-tree recalculations. The BPDUs sent out by the network attacker's system announce that the attacking system has a lower bridge priority. If successful, the network attacker can see a variety of frames. Figure 6 illustrates how a network attacker can use Spanning-Tree Protocol to change the topology of a network so that it appears that the network attacker's host is a root bridge with a higher priority. By transmitting spoofed Spanning-Tree Protocol packets, the network attacker causes the switches to initiate spanning-tree recalculations that then result in the two connections to the network attacker's system to forward packets.

Figure 6
Traffic Interception Using Spanning-Tree Protocol

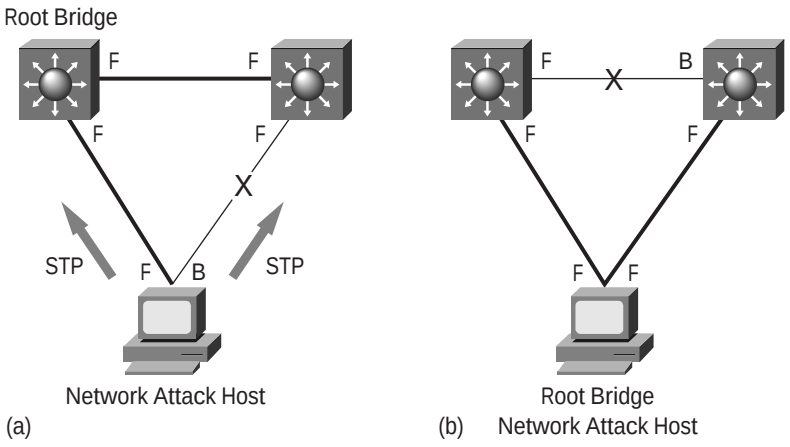

Network Attack Mitigation

To mitigate Spanning-Tree Protocol manipulation use the root guard and the bpdu guard enhancement commands to enforce the placement of the root bridge in the network as well as enforce the Spanning-Tree Protocol domain borders. The root guard feature is designed to provide a way to enforce the root-bridge placement in the network. The Spanning-Tree Protocol bpdu guard is designed to allow network designers to keep the active network topology predictable. While bpdu guard may seem unnecessary given that the administrator can set the bridge priority to zero, there is still no guarantee that it will be elected as the root bridge because there might be a bridge with priority zero and a lower bridge ID.

Command Samples to Mitigate Spanning-Tree Protocol Manipulation Attacks
CatOS> set spantree portfast bpdu guard enable
CatOS> set spantree guard root *mod_num/port_num*

MAC Spoofing Attack

MAC spoofing attacks involve the use of a known MAC address of another host to attempt to make the target switch forward frames destined for the remote host to the network attacker. By sending a single frame with the other host's source Ethernet address, the network attacker overwrites the CAM table entry so that the switch forwards packets destined for the host to the network attacker. Until the host sends traffic it will not receive any traffic. When the host sends out traffic, the CAM table entry is rewritten once more so that it moves back to the original port.

Figure 7 shows how MAC spoofing works. In frame (i) the switch has learned that Host A is on port 1, Host B is on port 2, and Host C is on port 3. Host B sends out a packet identifying itself as Host B's IP address but with Host A's MAC address or another packet with the same IP address and MAC address combination (ii). This traffic causes the frame to move the location of Host A in its CAM table from port 1 to port 2. Traffic from Host C destined to Host A is now visible to Host B (iii). To correct this situation, Host A must send out traffic on the switch port for the switch to "relearn" the location of Host A's MAC address (iv).

Figure 7
MAC Spoofing Attack

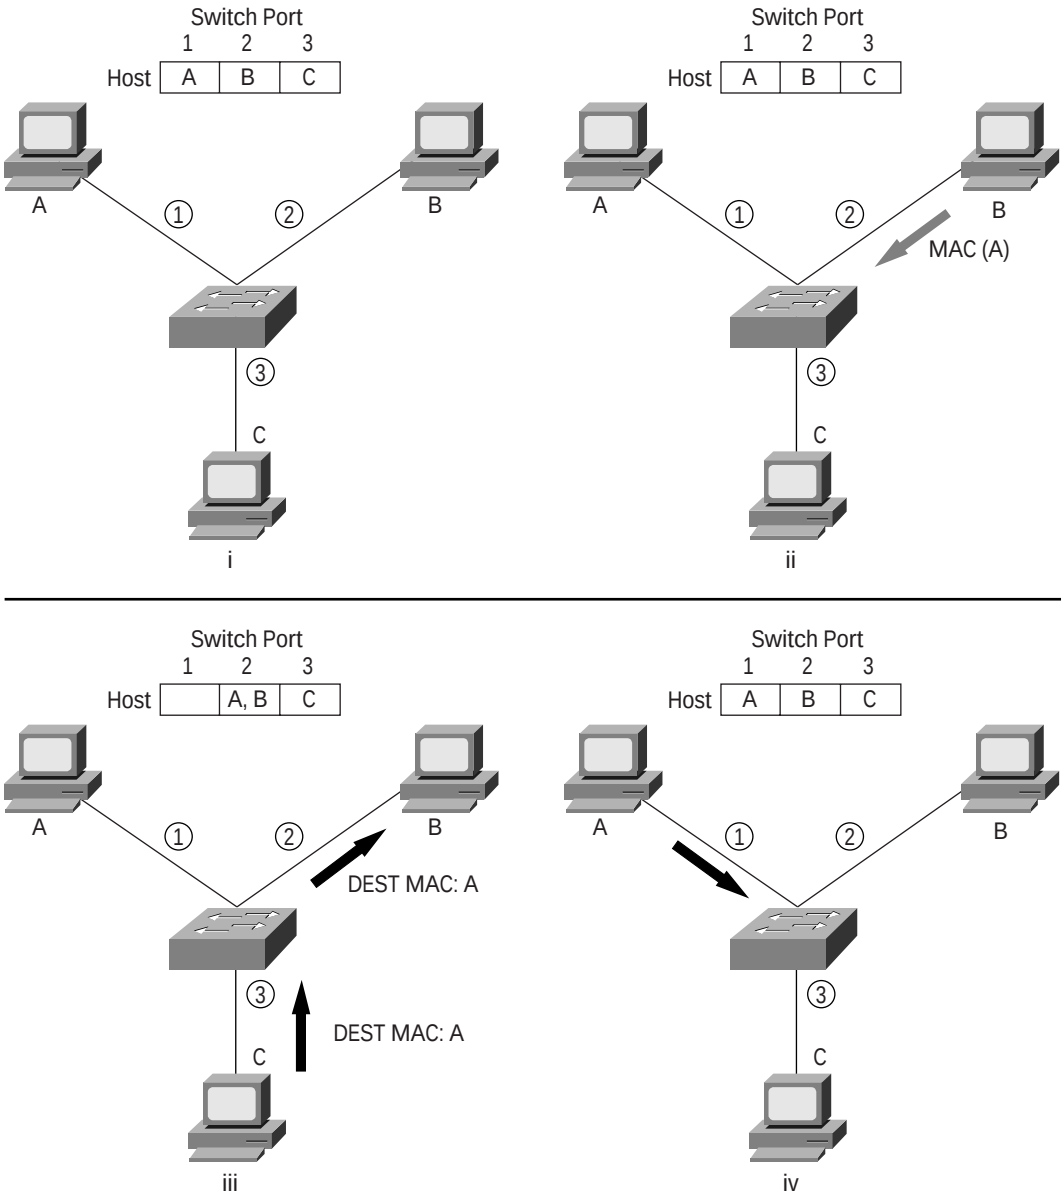

Address Resolution Protocol Spoofing—A variation of the MAC spoofing attack uses the Address Resolution Protocol (ARP) to insert entries into the CAM table. ARP is used to determine by a host to learn the MAC address of a destination host. Whereas the sender may know the IP address of the destination, it may not necessarily know the MAC address of the destination. To learn this information it constructs an ARP request packet asking the question "Who has IP address x.x.x.x? Tell me." The target responds with an ARP reply stating "x.x.x.x is at yy:yy:yy:yy:yy:yy." This ARP reply provides the sender with an IP-to-MAC mapping, which is then stored in the sender's ARP cache. The switch that the target is connected to also learns the MAC address of the destination host and what switch port connects to it. This information is then stored in the CAM table.

By crafting an ARP reply, a network attacker can make his or her system appear to be the destination host sought by the sender. The ARP reply causes the sender to store the MAC address of the network attacker's system in the ARP cache. This MAC address is also stored by the switch in its CAM table. In this way the network attacker has inserted the MAC address of his or her system into both the switch's CAM table and the sender's ARP cache. This allows the network attacker to intercept frames destined for the host that he or she is spoofing.

Network Attack Mitigation

Use the port security commands to mitigate MAC-spoofing attacks. The port security command provides the capability to specify the MAC address of the system connected to a particular port. The command also provides the ability to specify an action to take if a port-security violation occurs. However, as with the CAM table-overflow attack mitigation, specifying a MAC address on every port is an unmanageable solution. Hold-down timers in the interface configuration menu can be used to mitigate ARP spoofing attacks by setting the length of time an entry will stay in the ARP cache. However, hold-down timers by themselves are insufficient. Modification of the ARP cache expiration time on all end systems would be required as well as static ARP entries. Even in a small network this approach does not scale well. One solution would be to use private VLANs to help mitigate these network attacks.

Command Samples to Mitigate MAC Spoofing (Cisco Catalyst® Operating System [Catalyst OS])
CatOS> set port security *mod_num/port_num* **enable** *[mac_addr]*
CatOS> set port security *mod_num/port_num mac_addr*
CatOS> set port security *mod_num/port_num* **violation** *{shutdown | restrict}*
CatOS> show port *[mod_num[/port_num]*
Command Samples to Mitigate MAC Spoofing (Cisco IOS® Software)
IOS(config-if)# port security max-mac-count {1-132}
IOS(config-if)# port security action {shutdown|trap}
IOS(config-if) # arp timeout *seconds*

Private VLAN Attacks

While private VLANs are a common mechanism to restrict communications between systems on the same logical IP subnet, they are not a full-proof mechanism. Private VLANs work by limiting the ports within a VLAN that can communicate with other ports in the same VLAN. Isolated ports within a VLAN can communicate only with promiscuous ports. Community ports can communicate only with other members of the same community and promiscuous ports. Promiscuous ports can communicate with any port. One network attack capable of bypassing the network security of private VLANs involves the use of a proxy to bypass access restrictions to a private VLAN.

Proxy Attack—In this network attack against private VLANs, frames are forwarded to a host on the network connected to a promiscuous port such as a router. In Figure 8, the network attacker sends a packet with the source-IP and MAC address of his or her device, a destination IP address of the target system, but a destination MAC address of the router. The switch forwards the frame to the router's switch port. The router routes the traffic, rewrites the destination MAC address as that of the target, and sends the packet back out. Now the packet has the proper format as shown in Figure 8 and is forwarded to the target system. This network attack allows only for unidirectional traffic

because any attempt by the target to send traffic back will be blocked by the private VLAN configuration. If both hosts are compromised, static ARP entries could be used to allow bidirectional traffic. This scenario is not a private VLAN vulnerability because all the rules of private VLANs were enforced; however, the network security was bypassed.

Figure 8
Private VLAN Proxy Attack

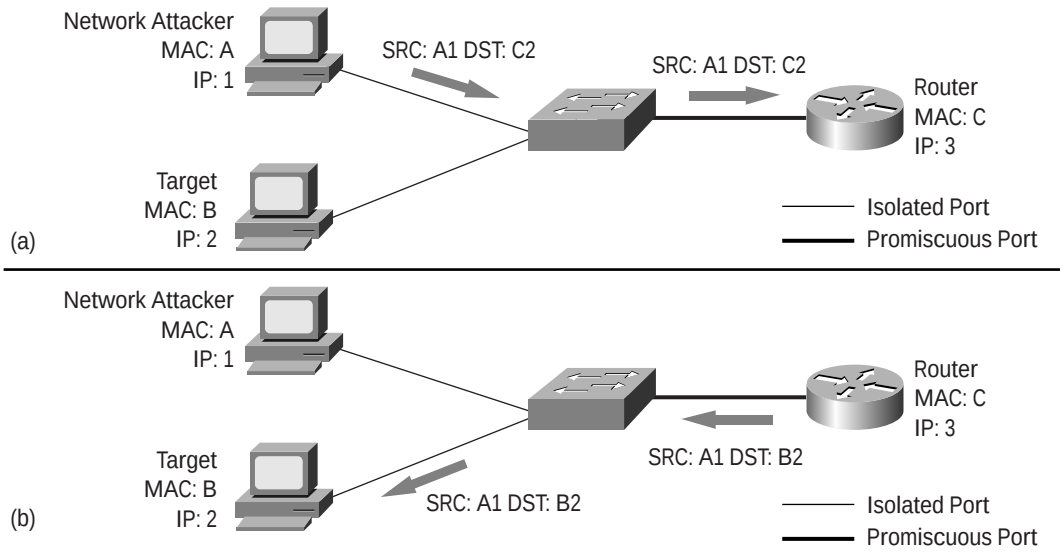

Network Attack Mitigation

Configure access control lists (ACLs) on the router port to mitigate private VLAN attacks. Virtual ACLs can also be used to help mitigate the effects of private VLAN attacks. An example of using ACLs on the router port is if a server-farm segment were 172.16.34.0/24, then configuring the following ACLs on the default gateway would mitigate the private VLAN proxy attack.

Command Samples to Mitigate Private VLAN Proxy Attack
IOS(config)# access-list 101 deny ip 172.16.34.0 0.0.0.255 172.16.34.0 0.0.0.255 log
IOS(config)# access-list 101 permit ip any any
IOS(config-if)# ip access-group 101 in

DHCP Starvation

A DHCP starvation attack works by broadcasting DHCP requests with spoofed MAC addresses. If enough requests are sent, the network attacker can exhaust the address space available to the DHCP servers for a period of time. This is a simple resource starvation attack just like a SYN flood is a starvation attack. The network attacker can then set up a rogue DHCP server on his or her system and respond to new DHCP requests from clients on the network. Exhausting all of the DHCP addresses is not *required* to introduce a rogue DHCP server, though. As stated in RFC 2131:

"The client collects DHCPOFFER messages over a period of time, selects one DHCPOFFER message from the (possibly many) incoming DHCPOFFER messages (for example, the first DHCPOFFER message or the DHCPOFFER message from the previously used server) and extracts the server address from the 'server identifier' option in the DHCPOFFER message. The time over which the client collects messages and the mechanism used to select one DHCPOFFER are implementation dependent."

By placing a rogue DHCP server in the network, a network attacker can provide clients with addresses and other network information. Since DHCP responses typically include default gateway and DNS server information, the network attacker can supply his or her own system as the default gateway and DNS server resulting in a "man-in-the-middle" attack.

Network Attack Mitigation

The techniques that mitigate CAM table flooding also mitigate DHCP starvation by limiting the number of MAC addresses on a switch port. As implementation of RFC 3118, Authentication for DHCP Messages, increases, DHCP starvation attacks will become more difficult. Finally, DHCP option 82, DHCP Relay Agent Information Option, helps to mitigate this network attack by enabling a DHCP relay agent to include information about itself when forwarding client-originated DHCP packets to a DHCP server. The server can use this information to implement IP address or other parameter-assignment policies.

Cisco Discovery Protocol

The Cisco Discovery Protocol runs at Layer 2 and allows Cisco devices to identify themselves to other Cisco devices. However, the information sent through Cisco Discovery Protocol is transmitted in cleartext and unauthenticated. Cisco Discovery Protocol is necessary for management applications and cannot be disabled without impairing some network-management applications. However, Cisco Discovery Protocol can be selectively disabled on interfaces where management is not being performed.

Network Attack Mitigation

Use Cisco Discovery Protocol only where appropriate.

Command Samples for Controlling Cisco Discovery Protocol
Global configuration:
IOS# no cdp run
Interface configuration menu:
IOS(config-if) no cdp enable

VLAN Trunking Protocol

The VLAN Trunking Protocol (VTP) is a Layer 2 messaging protocol that allows network administrators to centrally manage the addition, deletion, and renaming of VLANs. VTP is typically configured as a domain (also called a VLAN management domain) composed of one or more interconnected switches. The switches share the VTP-management domain name. Changes to the VTP domain can be made either through the command-line interface (CLI) or through Simple Network Management Protocol (SNMP) and are propagated to member switches through VTP advertisements. If a switch receives a VTP advertisement over a trunk link and it is not configured to be a *transparent* switch, it inherits the VTP domain name and configuration-revision number.

VTP security is provided through a password that is entered into the VTP database on all of the switches. This shared password is used to authenticate VTP advertisements.

At the present time no vulnerabilities have been identified or published with regard to VTP. It is theoretically possible to forge VTP packets and inject them into a VLAN management domain if the network attacker can configure his or her connection as a trunk link. In this way a network attacker could add or remove VLANs from the VTP domain as well as create Spanning-Tree Protocol loops. While the level of difficulty of such an attack remains high, the possibility is not inconceivable. VTP remains a valued method to centrally manage VLANs throughout an enterprise. It is recommended that a VTP password be set throughout the VTP domain to prevent the possibility of forging VTP advertisements.

Network Attack Mitigation

Assign a VTP password in the VTP management domain.

Command Samples for Assigning a VTP Domain Password
Global configuration: **CatOS> (enable) set vtp passwd** *passwd*

Summary

Although security attacks on networks aren't new events, attacks that use Layer 2 to bypass VLAN restrictions are quickly gaining sophistication and popularity. To mitigate the effects of these attacks as much as possible, the following precautions are recommended:

- Manage switches as securely as possible. Use Secure Shell (SSH) Protocol if possible, or an out-of-band management system. Avoid the use of cleartext management protocols such as Telnet or SNMP Version 1.
- Use IP-permit lists to restrict access to management ports.
- Selectively use SNMPv3 and treat community strings like root passwords.
- When SNMPv3 is used as a management protocol, restrict management access to the VLAN so that entities on untrusted networks cannot access management interfaces or protocols.
- Always use a dedicated VLAN ID for all trunk ports.
- Avoid using VLAN 1 for network management.
- Set all user ports to non-trunking mode.
- Deploy port security where possible for user ports. When feasible, configure each port to associate a limited number of MAC addresses (approximately two to three). This will mitigate MAC flooding and other network attacks.
- Have a plan for the ARP security issues in your network. Enable Spanning-Tree Protocol attack mitigation (BPDU Guard, Root Guard).
- Use private VLANs where appropriate to further divide Layer 2 networks.
- Use Cisco Discovery Protocol only where appropriate.
- Disable all unused ports and put them in an unused VLAN. This setup prevents network intruders from plugging into unused ports and communicating with the rest of the network.
- Use Cisco IOS Software ACLs on IP-forwarding devices to protect Layer 2 proxy on private VLANs.
- Eliminate native VLANs from 802.1q trunks.
- Use VTP passwords to authenticate VTP advertisements.
- Consider using Layer 2 port authentication such as 802.1X to authenticate clients attempting connectivity to a network.
- Procedures for change control and configuration analysis must be in place to ensure that changes result in a secure configuration. This is especially valuable in cases where several organizational groups may control the same switch, and even more valuable in network security deployments where even greater care must be taken.

By employing these additional precautions in the existing network configurations, the effects of Layer 2-based attacks can be significantly mitigated. Although not all of these precautions are necessary, depending on the network configuration, they do represent a single set of recommendations that can be used in addition to current network security guidelines.

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

RD/LW4135 01/03