



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Abstract

Ethernet Switching technology is increasing rapidly. As it does, it becomes much more important to include switches in a comprehensive security plan. Multilayer switching, Virtual Local Area Networks, or VLANs, VLAN trunking, flow-based switching, advanced functionality modules, such as intrusion detection, or network monitoring, and other enhancements and features are now commonplace on high-end switches and are rapidly migrating to lower-end switches. In some cases, multilayer switches are displacing routers. There are very real security implications that must be considered when including these technologies in a network design.

This paper focuses on the security considerations brought about when advanced switching technology is incorporated in a network. I have attempted to look at this subject with a broad enough focus so as to render it vendor and software version neutral. Therefore this document assumes reader familiarity with the commands to configure their switch, and does not give specific commands for switch configuration. First I will provide an overview of some of the key technologies appearing in switches. Next I will examine the security implications, both positive and negative, that these technologies present. Following this are some best practices to follow, both in evaluating the use, and in configuring advanced switching technologies. Finally I will list areas where I feel more research needs to be done in order to be able to securely incorporate advanced switching technologies in a network.

Considerations when deciding to use advanced Ethernet switching technologies

Much has been written about the use of switches in a security context. The consensus among many network security experts is that switches provide security risks which make them unsuitable for use in a secure environment. Rik Farrow wrote in his paper "VLAN Insecurity"¹ that switches were not designed to provide security. He writes that any time that you need to segregate networks for serious security purposes, you not use a switch. He makes a strong case that switches rely on software and configuration to provide any type of security and not physical isolation. This could also be said about routers and firewalls which rely on configuration to provide security. I feel that if care is taken with the configuration, and the hazards are known, there are places for Ethernet switching in an organization's security architecture.

New switches are being manufactured with advanced features that I feel warrant additional consideration for possible use in a secure environment. Multilayer switches are, in some cases, displacing routers; this complicates the issue. Now, instead of

¹ Farrow, Rik

being able to state simply that switches shouldn't be used in a security environment, we find that we have to integrate switches, and treat them as routers. In this paper I am proposing that since switches are ubiquitous in networks, we need to accommodate them. If the proper guidelines are followed, and switches aren't used as the primary means of providing security, they can be a valuable layer in creating defense in depth. There are several areas where switches can significantly add to the overall security of a network. One is through the use of inter-VLAN Access Control Lists, or ACLs; another is intrusion detection, and network monitoring.

I'm primarily going to focus on the Cisco Catalyst 6500 series of Ethernet switches, however I should add that other network equipment vendors are adding similar features to their products. At the end of this paper is a table which provides sources of further security information for many multilayer switch vendors.

Overview of advanced Ethernet switching technologies

I've included this next section for readers who would like background information on some of the more common advanced Ethernet switching technologies. If you are familiar with these technologies and want to know about the security implications, you may want to skip to the section titled "Security Implications of advanced Ethernet switching technologies."

VLANs

VLANs, or Virtual Local Area Networks, allow a switch to act as multiple switches. A single physical switch can be divided into multiple virtual switches. Each virtual switch becomes its own broadcast domain, or VLAN. The following drawing illustrates this. The rectangular box at the bottom of figure 1 represents a 6 port Ethernet switch.




Figure 1: 6 port switch with 3 VLANs

In figure 1, the switch has been configured with three VLANs creating three broadcast domains, which are color coded. Each VLAN exists at layer 2 of the ISO stack. The switch is designed so that traffic from one VLAN cannot be seen on another VLAN without going through a Layer 3 device. As we shall see, this is not always the case. By design, a yellow device in the drawing above can not connect to a blue or green device. The workstation attached to port 1 can only see traffic from ports 1, and 2.

VLAN Trunking

VLANs can also be extended across multiple switches. This is accomplished through the use of VLAN Trunking. To use VLAN trunking, a port on each switch is designated as a trunk port and is assigned a native VLAN. Next VLANs are added to the trunk. Additional configuration options can limit the VLANs that are carried on the trunk.

The IEEE standard, 802.1Q² which specifies VLAN trunking, provides for a 4 byte tag which is positioned directly after the source Media Access Control, or MAC address, of the frame. It should be noted that there are no provisions for authentication or validation of the data in the 802.1Q tag. The following drawing illustrates VLAN trunking.

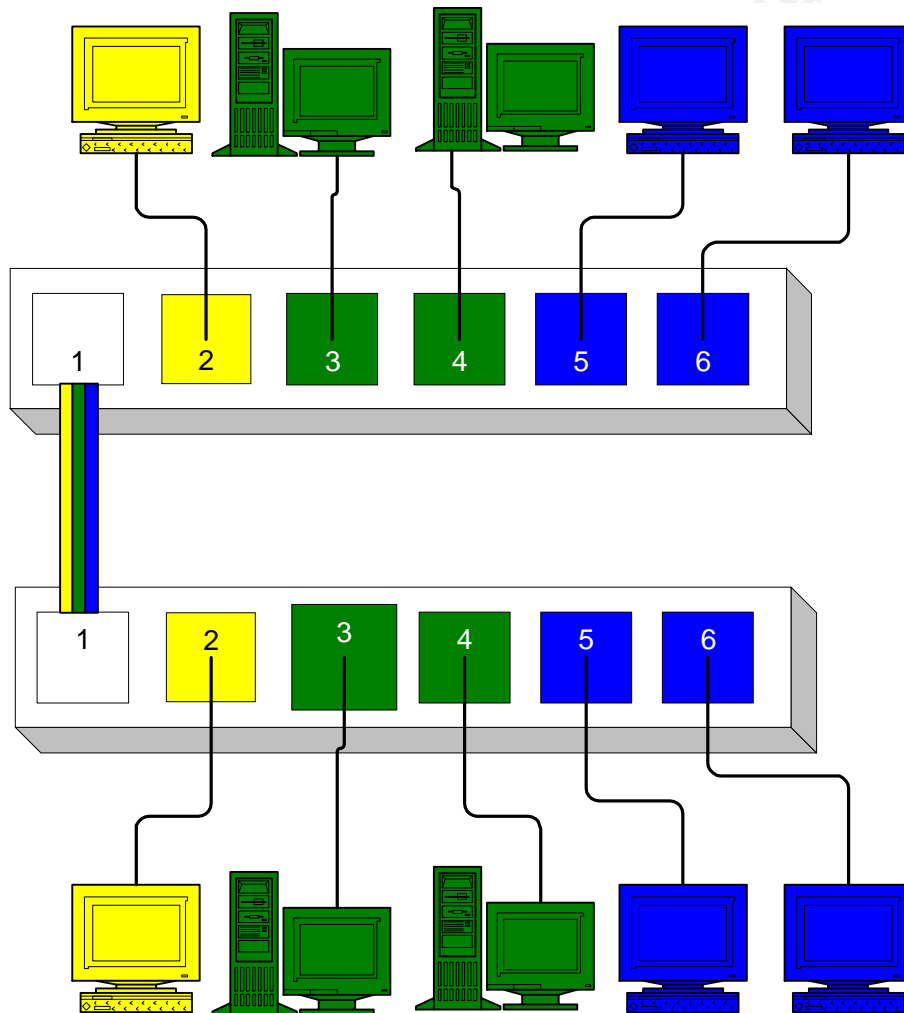


Figure 2: VLAN Trunking between 2 switches (3 VLANs)

Port 1 on both the upper and lower switches are configured as trunk ports. In figure 2 they are shown as carrying traffic from the yellow, green and blue VLANs. This doesn't always have to be the case. VLAN trunking protocols allow for a trunk to limit the

² IEEE Standards 802.1Q-1998

VLANs that it carries. This is a significant point that we will look at when we need to secure this network. With trunking configured, the VLANs now span the two switches. The workstation attached to port 4 on the upper switch can see traffic from ports 3 and 4 on the upper switch and ports 3, and 4 on the lower switch.

Multilayer switches.

In order to be able to move traffic from one VLAN to another, we need a Layer 3 device. This can be an external router as shown in figure 3.




Figure 3: 2 Switches with routing between 2 VLANs

A router is shown attached to port 3 on the upper switch and port 5 on the lower switch. This provides layer 3 connectivity between the blue VLAN and the green VLAN. Now

our workstation attached to port 4 on the upper switch can see all the traffic in both the green and blue VLANs. Note that the yellow VLAN is still isolated from the green and blue VLANs since it does not connect to a layer 3 device.

Another method to provide layer 3 connectivity between VLANs is through the use of a multilayer switch. This is becoming more common. Figure 4 illustrates multilayer switching.




Figure 4: All VLANs are connected through a multilayer switch

The upper switch is now shown as a multilayer switch. This means that it operates as a conventional switch, at layer 2, and as a router, at layer 3. Provided that all of the VLANs are connected at layer 3 in the upper switch, there is full connectivity among all of the devices in the drawing. This is because the lower switch is connected through a trunk, which carries all three VLANs.

Flow switching.

There is a significant performance benefit possible with multilayer switching. Cisco accomplishes this through the use of flow-based switching. Cisco considers a flow to

be the traffic between the same source IP address and port, and destination IP address and port. Using the last illustration, let's look at how flow-based switching works.

If the workstation attached to port 6 of the upper switch wants to connect to the server connected to port 3 of the upper switch, it must go through a layer 3 device since each device is in a different VLAN. If it went through a router, every packet would need to be examined at layer 3 and a routing decision would need to be made. Because the layer 3 routing engine is closely integrated with the layer 2 portion of the switch, it can make the routing decision once, and then by recording the flow, and making that information available to the layer 2 switching engine, pass the rest of the traffic through only the layer 2 switching engine. Only the initial packets of the flow would have to go through the layer 3 routing engine. The rest of the flow would be handled much like 2 ports in the same VLAN.

Inter-VLAN ACLs

Utilizing a multilayer switch that allows for inter-VLAN ACLs can add significantly to the security of a network. It allows some of the access controls that are common in firewalls throughout the entire network. It would generally be cost and performance prohibitive to firewall every subnet on a network; however due to the performance and cost benefits of multilayer switching, it now become much more feasible to subject all traffic in the network to ACLs. By moving inter-VLAN traffic to a layer 3 engine, we can apply ACLs and make routing decisions about the traffic, however once that is done, through the use of flows, we can pass the rest of the traffic without the overhead of a layer 3 device.

Intrusion Detection Modules

In a network that utilizes multilayer switching and VLAN trunking, it is possible to bring all traffic through a single device. This creates a very convenient place to do intrusion detection. Cisco does this with its Intrusion Detection Module for the 6500 series of switches. This module acts similarly to a VLAN trunk, in that it has a presence on each VLAN. This allows for a single probe to monitor multiple VLANs. Because the connection is through the backplane, it is able to operate with high traffic loads.

Security Implications of advanced Ethernet switching technologies

Ethernet switching takes place primarily at layer 2, with multilayer switching involving layer 3. Layer 2 presents challenges in securing a network since it is often overlooked.³ Too often, the network security staff focuses on layer 3 and above, but as we will see, security vulnerabilities at layer 2 can and do affect everything that lies above it. Another potential problem with layer 2 devices is that they usually ship with security wide open. One notable exception at the time of this writing is switches manufactured by Alcatel, which ship locked down by default.⁴ The overall state of information security would be greatly enhanced if more vendors, both of hardware and software, adopted this stance.

³ Howard, Connie

⁴ Alcatel Technology Brief

Ethernet switching also presents opportunities for increasing security, as we shall explore later. First let's look at specific vulnerabilities that we may see which are specific to Ethernet switches. It is beyond the scope of this paper to cover every vulnerability so I have limited it to the most common: MAC Flooding, MAC Spoofing and VLAN Hopping.

MAC Flooding

An Ethernet switch works much like a bridge. In fact it really is just a multiport bridge. It connects devices on the same subnet, or broadcast domain. Like a bridge, it looks at the MAC address of the packets that enter each interface and it records the source MAC address in the Content-Addressable Memory (CAM) lookup table in its memory.

When a switch receives a frame on a port it first looks at the source of the frame. If the source is not known, it updates the CAM table, noting the MAC address and the port on which it entered. Next it looks at the destination MAC address, and checks to see if that MAC address exists in its CAM table. If it does, it copies the frame to the port associated with the MAC address. If it doesn't have an entry for the MAC address in the CAM table, it will flood, or copy, the frame out of every port. In this mode, it acts just like a hub. If a response comes back, then the switch updates its CAM table with the MAC address and port of the responding device. A subsequent packet could then be copied to the proper port instead of being flooded out of all ports. Over time the number of entries in the CAM table can grow, and since the memory is a finite resource, at some point, the number of entries may exceed the amount of memory allocated for this table.

One method which is used to keep the CAM table from growing too large is called aging. If the switch doesn't hear from a specific device before its aging timer expires, it removes the entry from the table. Generally, this is set to about 5 minutes. If it is set too high, the CAM table can fill up more easily. If it is set too low, the switch may needlessly flood frames for devices that it knew about.

What would happen if the entries were added to the CAM table at a rate faster than the aging timer removed them? The CAM table would fill up and the switch would not be able to add a new MAC address to the table. Let's look at this on a hypothetical 4 port switch which only holds 16 MAC addresses in its CAM table.

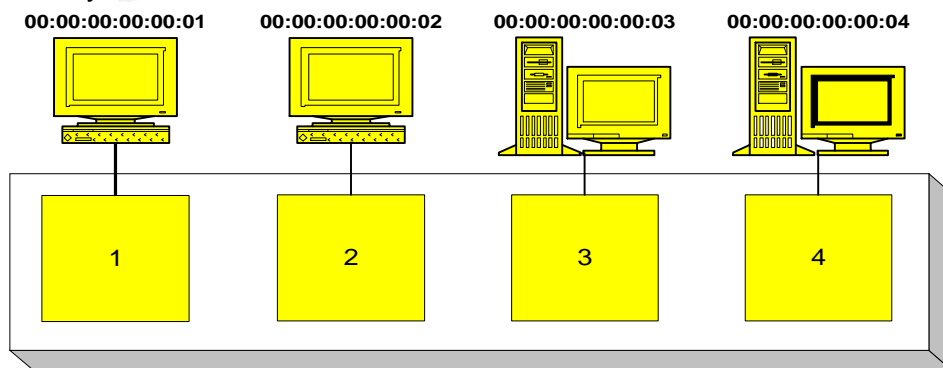


Figure 5: MAC addresses of devices attached to a switch

Figure 5 shows this switch with 4 devices attached. If all 4 devices were on, the CAM table would be as shown in the first table.

Entry	Port	MAC Address
1	1	00:00:00:00:00:01
2	2	00:00:00:00:00:02
3	3	00:00:00:00:00:03
4	4	00:00:00:00:00:04
5		
6		
7		
8		
9		
A		
B		
C		
D		
E		
F		
10		

Entry	Port	MAC Address
1	1	00:00:00:00:00:01
2	2	00:00:00:00:00:02
3	3	00:00:00:00:00:03
4	4	00:00:00:00:00:04
5	4	00:00:00:11:11:10
6	4	00:00:00:11:11:11
7	4	00:00:00:11:11:12
8	4	00:00:00:11:11:13
9	4	00:00:00:11:11:14
A	4	00:00:00:11:11:15
B	4	00:00:00:11:11:16
C	4	00:00:00:11:11:17
D	4	00:00:00:11:11:18
E	4	00:00:00:11:11:19
F	4	00:00:00:11:11:1A
10	4	00:00:00:11:11:1B

If the device on port 4 was to start sending packets with forged MAC addresses, the CAM table would quickly fill as shown in the second table. At this point, the switch would not be able to add any more entries. The normal behavior for a switch under these conditions would be to flood every frame received to every port. In effect, the switch would become a hub.

A common belief is that a switch prevents the use of a sniffer to snoop LAN traffic. If a switch could be made to act like a hub, then sniffing would be possible. Refer back to figure 5. If the device attached to port 1 was trying to communicate securely with the device connected to port 3, traffic between the two would not normally be able to be seen on port 4. Now let's say that a malicious user had control of the device on port 4. By flooding port 4 with forged MAC addresses, he or she could cause the switch to act like a hub. It would then be very easy for the user on port 4 to sniff the traffic between ports 1 and 3 since all traffic would be flooded out of every port.

In reality switches have room for much more than 16 entries in the CAM table. It varies according to the model of the switch and in some cases the memory installed, but it is not unusual for there to be room for many thousands of entries. This doesn't provide a greater level of security, however. There is software readily available that can generate thousands of random forged packets per second. The Dsniff package which contains the program macof is the most common. Macof is capable of generating 155,000 MAC addresses on a switch per minute.⁵ A web search for "Dsniff" will turn up many responses, including sources of the software.⁶ Two very good sources of information can be found in the SANS InfoSec Reading Room. Lora Danielle covers the use and

⁵ Convery, Sean. Page 18

⁶ Song, Dug

installation of Dsniff in her paper, "Introduction to Dsniff".⁷ Of more importance to a security practioner is knowing if Dsniff is being used on a network. Richard Duffy's paper, "Finding Dsniff on your network"⁸ covers this subject in depth. **WARNING:** Running dsniff can have unpredictable effects on network switches. Be sure to get written permission before using dsniff. It is also advisable not to run dsniff in a production environment.

Protecting against MAC Flooding Attacks

The primary method to protect against MAC Flooding attacks is by configuring port security. There are several ways to do this. One method is to restrict a port to a specific IP address. At first glance it would appear that this has the additional advantage of preventing unknown devices from connecting to the switch; however MAC spoofing negates this benefit. It will however prevent someone from extending the network by connecting hubs or switches to a port. If you choose to restrict the port to a specific address, it can be learned from the first device that connects to the port, or it can be specified directly. A caveat is that there is a good deal of overhead, both management and performance, when you use this type of port security..

Another method is to limit the number of MAC addresses which can be associated with a single port. As long as the total number of addresses allowed for all ports doesn't exceed the size of the CAM table, the CAM table should never become filled. A port limited in this manner can be configured to shut down if the limit is reached, or it can reject additional addresses. Both methods have their tradeoffs. If you have the port shut down, you are more likely to know if the port is being used for malicious purposes. However you create a situation where someone using a tool like dsniff could do a DoS attack instead. A combination of methods is probably best. For switch ports used to connect to servers and routers, restrict the port to a specific address, and ignore any other addresses. For ports used to connect a single user device, having the port shut down, if more than a specified number of addresses are seen on it, is probably more secure.

MAC Spoofing

A MAC spoofing attack involves forging a packet with a MAC Address which is known to the switch. Let's take 4 network devices connected to a switch, as shown in the next illustration.

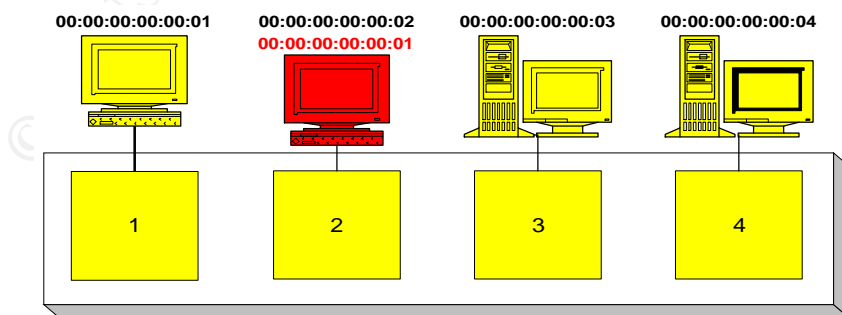


Figure 6: MAC spoofing illustration

⁷ Danielle, Lora

⁸ Duffy, Richard

The switch's CAM table associates 00:00:00:00:00:01 with port 1, 00:00:00:00:00:02 with port 2, 00:00:00:00:00:03 with port 3, and 00:00:00:00:00:04 with port 4. If an attacker on port 2 wanted to intercept the traffic between port 1 and 3, the attacker could send a packet with the address of 00:00:00:00:00:01. The switch would then update its CAM table to reflect that 00:00:00:00:00:01 was now connected to port 2. Traffic from the server connected to port 3 destined for the workstation on port 1 would now be redirected to the workstation on port 2.

As soon as the workstation on port 1 sent another packet to the switch, the switch would again update its CAM table to reflect the correct location of 00:00:00:00:00:01, and the switch would cease to send traffic directed for port 1 to port 2. This type of attack would be more effective if there was a simultaneous DoS attack to the workstation on port 1 which would keep that workstation from being able to send packets to the switch. This could be launched from the same host, on port 2, or another host, such as the one on port 4.

Protecting against MAC spoofing attacks.

The steps to protect against MAC spoofing attacks are primarily the same as those used to prevent against MAC flooding attacks. The most effective solution is to restrict a switch port to a single MAC address. It should also be noted that MAC spoofing and MAC flooding attacks are confined to a subnet. That is, they do not cross layer 3 boundaries through a router. This is not to say that an attacker cannot remotely assume control of a system from another subnet and then use that system to launch one of these attacks. This further points to the need for defense in depth.

VLAN Hopping

Dave Taylor and Steve Schupp reported in 1999 that under certain circumstances they were able to get a frame to jump from one VLAN to another without the intervention of a layer 3 device.⁹ Dave Taylor expanded upon his earlier research in his SANS paper in 2000.¹⁰ This paper is an excellent reference of the exact method used to get frames to hop VLANs. He found that when a frame was sent from a host in one VLAN to a host in another VLAN, the frame would be successfully delivered if the frame contained the VLAN ID for the second VLAN in the 802.1Q tag of the frame. This only worked when the two hosts were connected to different switches, connected with an 802.1Q trunk.

This pointed directly at a security problem in the way 802.1Q trunking takes place. It was also noted that when a frame was able to successfully hop VLANs, the source VLAN was the same as the native VLAN of the trunk port.

Cisco reports another method in which frames can hop VLANs, which it refers to as double tagging. This is documented in their SAFE Layer 2 whitepaper.¹¹ A double tagging attack involves frames which have two 802.1Q tags in the frame header. The first switch which sees the frame will strip the tag, and forward the frame. The resulting frame is then sent out with the second tag intact. If this frame is received on another

⁹ Taylor, Dave. Schupp, Steve

¹⁰ Taylor, Dave

¹¹ Cisco Systems, [SAFE Enterprise Layer 2 Addendum](#) Page 5

switch, that switch will forward the frame based on the information in the second 802.1Q tag.

There are serious consequences to VLAN hopping, especially if a switch is being used to segregate traffic for security purposes. The next diagram illustrates this.

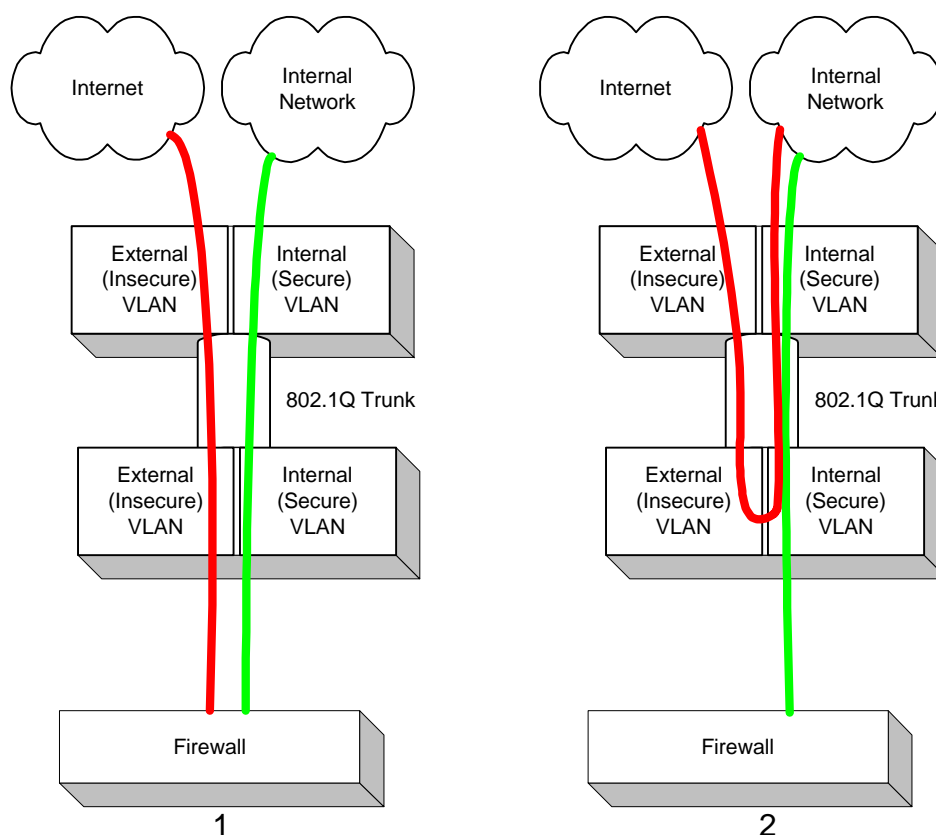


Figure 7: Problems with VLAN Hopping

In illustration number 1, we see two switches, and two VLANs utilizing 802.1Q trunking. The VLANs are being used to segregate secure and insecure traffic. Traffic from the internet passes from the first switch, over a trunk, into a second switch and into a firewall. At the firewall, rules are applied which deny or allow the traffic based on the organization's security policy. If the traffic is allowed, it passes through the switch, over the trunk, into the other switch, and into the organization's internal network.

In the second illustration we see the implications of VLAN hopping. Traffic from the internet passes into the first switch, over the trunk, and then hops to the internal VLAN, and enters the internal network. In this scenario traffic from the Internet is allowed to pass into the internal network without being subject to rule checking at the firewall. The organization's security policy regarding access from the internet is effectively negated.

Protecting against VLAN hopping attacks

Taylor's research provides valuable information on how to protect against VLAN hopping attacks. Recall that there were several conditions that had to be met before a frame would hop VLANs:

- The frame had to be tagged with the VLAN ID of the destination device
- The source VLAN had to be the same as the native VLAN of the trunk port
- In order for the frame to get an 802.1Q tag, it had to originate on a port that allowed trunking.
- The traffic had to utilize an 802.1Q trunk. That is, the traffic had to cross two switches.

We can prevent VLAN hopping by configuring the switch as follows:

- Insure that there is a unique native VLAN number for trunking ports. This will prevent the source VLAN from being the same as the native VLAN on the trunk port.
- Insure that trunking is explicitly disabled on all ports except those that are being used for trunking.
- Design the network architecture so that insecure traffic doesn't cross multiple switches. Figure 8 gives an example of this.




Figure 8: Architecture which defends against VLAN hopping

Opportunities to enhance security with Ethernet switching

As we have just seen, there are some potentially serious security implications when a switch is used to segregate secure and insecure network traffic. Why then would we want to consider using a switch in such a manner? I feel that there are two areas where if best practices are followed, a switch can add to the overall security of a network. They are through the use of Inter-VLAN ACLs, and through the use of Network Monitoring or Intrusion Detection Modules.

Inter-VLAN ACLs

A multilayer switch provides an excellent way to extend firewalling throughout the network through the use of Inter-VLAN ACLs. Currently it is often cost and performance prohibitive to apply firewall technology throughout an organization. Consider the following organization with four divisions.

One division is Accounting. This includes accounts receivable, accounts payable and payroll. They have 6 users and 2 servers for their accounting applications. The second division is Sales. They have 10 people and a single server which contains information used by the sales staff. The third division is Manufacturing which has 20 employees and 2 servers used to automate the manufacturing process. Lastly, we have Research, which consists of 6 individuals and 1 server. In an organization of this size it is not likely that firewalls would be installed between each division. Further, it is also not likely that it would be cost effective to install a router with the power and port density to efficiently handle ACLs to segment the traffic.

This is where a multilayer switch with Inter-VLAN ACLs can help. Due to the method that a multilayer switch uses flows, it can apply routing decisions and ACLs once for an entire flow. Each packet is no longer subject to rule checking as would be the case if ACLs were handled in a router. Now we can create VLANs for groups of systems and create a matrix of the permissions between each group. Using our fictional organization, we come up with the following:








































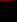

















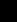







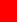
Access Permitted  Access Denied 								
	Accounting Users	Sales Users	Manufacturing Users	R&D Users	Accounting Servers	Sales Servers	Manufacturing Servers	R&D Servers
Accounting users								
Sales users								
Manufacturing users								
R&D Users								
Accounting Servers								
Sales Servers								
Manufacturing Servers								
R&D Servers								

Figure 9: Inter-VLAN Permission matrix

Looking at the first row of the matrix, we see that accounting users have access to the accounting servers and the sales servers. They are denied access everywhere else. Based on the matrix created, ACLs can be added to the multi-layer switching engine to enforce the organization's security policy in the switch, and reinforce the least privilege principle.

Network Monitoring and Intrusion Detection Modules.

Many vendors of Multi-Layer Ethernet switches are now offering an option to include Network monitoring and/or Intrusion Detection. I feel that a switch offers an excellent place to institute network based intrusion detection for several reasons:

- If all VLANs come through a single switch, it allows one place to look at network traffic, and eliminates the need to have many IDS probes.
- In some switch vendor's implementations of IDS, the probe listens on a VLAN without taking an IP address. In this manner the probe is not readily apparent to an attacker.
- Since logging of network traffic on multiple subnets is performed by a single device, logs will be consolidated and synchronized, which should make analysis easier
- Switch based intrusion detection modules can monitor traffic on the switch's backplane, theoretically offering higher performance.
- A single IDS monitoring both internal and external traffic is theoretically better able to correlate attack data.

If you consider the network for the fictional organization used in the inter-VLAN ACL example, there are 8 subnets in an organization of approximately 50 people. A switch based Network Intrusion Detection System, or NIDS could cover all 8 subnets with a single device, instead of 8 probes with traditional NIDS. Figure 10 shows how a switched base NIDS is able to monitor both internal and external traffic.




Figure 10: Switched based Network Intrusion Detection System

Best Practices when implementing advanced Ethernet switching technologies

It is apparent from the information above, that while there are security issues involved with using switches to handle traffic of varying security levels, there are also advantages. The challenge is to be able to utilize the advantages and negate the disadvantages. If this challenge is met, switches that provide advanced Ethernet switching technologies can be a valuable addition to an organization's defense in depth. The following list of best practices will allow you to implement high-end switches in a manner that enhances the overall security of your organization:

Layer 2 security in general:

- Create a policy that defines the security posture of Ethernet switches. Decide if you will allow mixed security traffic on the same switch. Define whether or not sensitive traffic is allowed on a VLAN at all
- Document your switch configurations. Include VLAN configuration, Trunk Configuration, Port Security. **Make sure that you know all the devices that your sensitive traffic may be carried on.**

Device Security:

- Only allow administrative access to your switch via secure methods, i.e. console port, SSH, SSL if using HTTP administration, VPN. Consider using strong authentication, such as S/Key or SecurID. Use IP permit lists to restrict access to management ports.
- Put the management interface into a dedicated VLAN (NOT VLAN 1) that is used just for management.
- Change all default passwords.
- Treat your switches like a router. Give them the same security considerations as you would a router. The Center for Internet Security has an excellent benchmark for securing Cisco routers.¹² Many of the steps in the benchmark apply to switches from any vendor.
- Insure that there is adequate physical security for your switches. Switches, unlike routers, are often placed in wiring closets that do not have the same level of physical security as the computer rooms that routers are often found in.
- Disable unused ports.
- Turn off Cisco Discovery Protocol, or CDP (Cisco only) on ports that do not connect to another Cisco device.
- Keep your switch code current. Subscribe to a notification service or mail list so that you are aware of bugs and vulnerabilities.
- If you use your switch to carry security sensitive information, then your switch becomes part of your security perimeter. **DON'T DO THIS UNLESS THE SWITCH IS UNDER THE CONTROL OF THE SECURITY OPERATIONS STAFF.**

¹² Center for Internet Security

Mitigating MAC Flooding and Spoofing Attacks:

- Use port security on all switch ports. Limit the number of CAM entries for each port or restrict the port to a specific MAC address.

Mitigating VLAN Hopping Attacks

- Do not trunk security sensitive traffic. Research indicates that VLAN hopping only takes place when a trunk is involved.¹³
- Create a separate VLAN used only as the native VLAN for trunk ports.
- Do not use VLAN 1 for anything. Especially trunking. Put unused ports in a VLAN created specifically for them, and do not allow that VLAN to be trunked.
- Explicitly disable trunking on all ports, and then turn on for ports which will participate in trunking. Do not allow ports to remain in the auto trunking mode.
- Prevent trunks from carrying security sensitive VLAN traffic
- Double tagging has a readily identifiable signature.¹⁴ Make sure that your Network Intrusion Detection System knows how to detect it.
- Do not use 802.1Q tags for any type of security decision since they are not authenticated.

Inter-VLAN ACLs

- Utilize Inter-VLAN ACLs to provide segmentation that enforces the least privileges concept.
- Consider not using flow-based switching for sensitive traffic until more research has been done. (see “Work to be done”, below)

Intrusion Detection/Network Monitoring Modules.

- Utilize available intrusion detection or network monitoring modules to provide IDS capabilities on all subnets.

Work to be done

- This paper limited its scope to the most common attacks. It does not cover all the attacks possible against a layer two switch. It should be extended to cover other attacks such as ARP spoofing, DHCP starvation, VLAN Trunking Protocol Attacks, Spanning Tree Protocol Attacks, Private VLANs, and VoIP Attacks. A good starting point is Shawn Convery’s presentation, “Hacking Layer 2: Fun with Cisco Switches”.¹⁵
- Create a “cookbook” of configuration commands, based on the best practices above, for the most common switches and software versions.
- Switches should ship in a secure state much like routers do. They should not allow communications unless enabled.
- Switches should defend against MAC Flooding attacks. It would be beneficial if switches were designed to fail-safe instead of fail-open. A switch with a CAM table overflow should not flood packets to every port.

¹³ Taylor, Dave

¹⁴ Convery, Sean pg. 28

¹⁵ Convery, Sean

- Is flow-based switching secure? This needs to be researched. I couldn't find any references which look at this. Since only the first packet in flow based switching is subject to routing decisions and ACL checking, I can see the potential to base an attack on flow-based switching. In such an attack, after the flow was established, and ACL checking done, malicious traffic could then be sent through the flow.

Multi Layer Switches Information Resources

The following table provides sources of additional security information by various switch vendors. The vendor list is from a spreadsheet from the NWFusion web site.¹⁶

Vendor	URL for Switch Security information/	URL for VLAN Hopping information
3Com	http://www.3com.com/corpinfo/en_US/pressbox/press_release.jsp?INFO_ID=139871	
Alcatel	http://www.ind.alcatel.com/library/techbrief/TB_Security-By-Default_2H-02.pdf	
Avaya	http://www1.avaya.com/enterprise/whitepapers/msn1841.pdf http://www1.avaya.com/enterprise/whitepapers/vlan-tutorial.pdf	
Cisco	http://www.cisco.com/en/US/netsol/ns110/ns170/ns171/ns128/networking_solutions_white_paper09186a008014870f.shtml	http://www.cisco.com/en/US/products/hw/switches/ps708/products_white_paper09186a008013159f.shtml http://www.cisco.com/en/US/about/ac123/ac114/ac173/ac222/about_cisco_packet_feature09186a0080142deb.html http://www.cisco.com/en/US/netsol/ns110/ns221/ns223/ns227/networking_solutions_white_paper09186a00800a1195.shtml
Extreme	http://www.extremenetworks.com/libraries/whitepapers/technology/Security.asp http://www.extremenetworks.com/common/asp/frameHandler.asp?go=/libraries/casestudies/Security_SB.pdf	http://www.extremenetworks.com/libraries/techbriefs/Metro_TG_vMAN.asp
Foundry	http://www.foundrynet.com/solutions/appNotes/ironShieldSecurity.html	
Nortel	http://www.nortelnetworks.com/corporate/events/2001d/security_eseminar/collateral/55065_25_11_01.pdf	

¹⁶ Network World Fusion

References

Alcatel Technology Brief. "Security by Default."

URL: http://www.ind.alcatel.com/library/techbrief/TB_Security-By-Default_2H-02.pdf (8 May 2003).

Bourke, Tony. "VLAN with a Plan." HostingTech, Issue 1.8

URL: http://www.hostingtech.com/nm/01_08_vlan.html (9 May 2003).

Center for Internet Security. "CIS Level-1/Level-2 Benchmark and Audit Tool for Cisco IOS Routers." March, 2003.

URL: http://www.cisecurity.org/bench_cisco.html (13 May 2003).

Cisco Systems. Building Cisco Multilayer Switched Networks, Student Guide, Revision 1.0a. 2000.

Cisco Systems. "SAFE Enterprise Layer 2 Addendum." SAFE Blueprint.

URL: http://www.cisco.com/warp/public/cc/so/cuso/epso/sqfr/sfblu_wp.pdf (8 May 2003).

Cisco Systems. "VLAN Security White Paper."

URL: http://www.cisco.com/warp/public/cc/pd/si/casi/ca6000/prodlit/vlnwp_wp.pdf (9 May 2003).

Clark, Kennedy. Hamilton, Kevin. Cisco LAN Switching. Indianapolis, IN: Cisco Press, 1999.

Convery, Sean. "Hacking Layer 2: Fun with Cisco Switches." 2002.

URL: <http://www.blackhat.com/presentations/bh-usa-02/bh-us-02-convery-switches.pdf> (13 May 2003).

Danielle, Lora. "Introduction to Dsniff." SANS InfoSec Reading Room. June 1, 2001.

URL: <http://www.sans.org/rr/audit/dsniff.php> (8 May 2003).

Duffy, Richard. "Finding Dsniff on Your Network." SANS InfoSec Reading Room.

November 28, 2001. URL: <http://www.sans.org/rr/penetration/dsniff.php> (8 May 2003).

Farrow, Rik. "VLAN Insecurity."

URL: <http://www.spirit.com/Network/net0103.html> (30 April 2003).

Howard, Connie. "Layer 2: The Weakest Link." Packet, Vol. 15 Num. 1. January, 2003

URL: http://www.cisco.com/en/US/about/ac123/ac114/ac173/ac222/about_cisco_packet_feature09186a0080142deb.html (8 May 2003).

“IEEE Standards 802.1Q-1998. Virtual Bridged Local Area Networks.” IEEE Standards for Local and Metropolitan Area Networks. December 8, 1998.
URL: <http://standards.ieee.org/getieee802/download/802.1Q-1998.pdf> (8 May 2003).

Network World Fusion. “Buyers Guide: Ethernet Switches. Buyers Guide Chart.” Network World Magazine. August 26, 2002.
URL: <http://www.nwfusion.com/downloads/0826BGSwitches.xls> (8 May 2003).

Rossi, Louis R. Rossi, Louis D. Rossi, Thomas L. Cisco Catalyst LAN Switching. New York, NY: McGraw-Hill, 2000.

Pollino, David. Schiffman, Mike. “Secure Use of VLANs: An @stake Security Assessment”. August, 2000.
URL: http://www.cisco.com/warp/public/cc/pd/si/casi/ca6000/tech/stake_wp.pdf (9 May 2003).

Song, Dug. “dsniff.” URL: <http://monkey.org/~dugsong/dsniff/> (9 May 2003).

Taylor, David. Schupp, Steve. “VLAN Security.” BugTraq. Sep 1 1999.
URL: <http://www.securityfocus.com/archive/1/26008> (30 April 2003).

Taylor, David. “Are there Vulnerabilities in VLAN implementations?” July 12, 2000.
URL: <http://www.sans.org/resources/idfaq/vlan.php> (9 May 2003).

Turner, Arron D. “Network Insecurity with Switches” August 29, 2000.
URL: http://synfin.net/docs/switch_security.html (9 May 2003).