Application Note

Virtual LAN Security Best Practices

Independent security research firm @stake [9] recently conducted a Security Review [1] of the virtual LAN (VLAN) technology on the Cisco Catalyst 2950, Catalyst 3550, Catalyst 4500, and Catalyst 6500 series switches. Although no intrinsic security weaknesses emerged from this review, it has been pointed out that an improper or inadequate switch configuration can be the source of undesired behavior and possible security breaches.

Over the past years, Cisco Systems has been advocating best-practices guidelines for secure network configuration in several documents. The *SAFE Blueprint* [2] or the *Best Practices for Catalyst 4500, 5000,* and *6500 Series Switches* [3] are examples of such documents. However, there has been no single document that collects all of the VLAN-related best practices for easier perusal by customers and field engineers.

The purpose of this paper is to present in a comprehensive way all of the recommendations that Cisco engineers have accumulated to aid with the proper configuration of VLANs on Cisco switches. At the same time, through direct-to-the-point descriptions, the main results of the @stake testing will be explained and the security threats demystified.

Basic Security

Any attempt to create a secure switched network starts from basic security principles. And in particular, basic rules such as the ones highlighted in the *SAFE* best practices [2] are the cornerstone of any design of secure switched networks.

If a user does not want one of his or her devices to be tampered with, physical access to the device must be strictly controlled. Furthermore, it is important for any network administrator to use all the proven security tools available on Cisco platforms: from the very basic configuration of system passwords, the use of IP permit filters, and login banners, all the way to more advanced tools such as RADIUS, TACACS+, Kerberos, SSH, SNMPv3, IDS, and so forth. (More details are provided in [3].)

Only after all the basic security components are in place, is it possible to turn attention to more sophisticated security details. In the following sections, VLAN-related issues will be explained.

Virtual LANs

A Layer 2 (L2) switch is a device capable of grouping subsets of its ports into virtual broadcast domains isolated from each other. These domains are commonly known as virtual LANs (VLANs).

The concept of VLAN is akin to other concepts in the networking world where traffic is identified by the use of a tag or

label. Identification is crucial for a L2 device to be able to isolate ports and properly forward the traffic received. As we will see later, lack of identification is sometimes a cause of insecurity and needs to be avoided.

If any packet in a device is tightly coupled to an appropriate VLAN tag, it is always possible to reliably discriminate traffic into separate and independent domains. This is the basic premise of VLAN-based switching architectures.

In particular, Cisco devices work in accordance with popular VLAN tagging technologies like ISL or 802.1Q across physical links (sometimes referred to as *trunks*) and employ advanced tagging techniques to preserve the VLAN information internally and use it for the purpose of traffic forwarding.

The simple observation that can be made at this point is that if a packet's VLAN identification cannot be altered after transmission from its source and is consistently preserved from end to end, then VLAN-based security is no less reliable than physical security.

Further on this topic in the following sections.

Control Plane

Malicious users often seek to gain access to the management console of a networking device, because if they are successful they can easily alter the network configuration to their advantage.

In a VLAN-based switch, in addition to having a direct connection to an out-of-band port, the management CPU can use one or more VLANs for in-band management purposes. It also uses one or more VLANs to exchange protocol traffic with other networking devices.

As basic physical security guidelines require networking equipment to be in a controlled (locked) space, VLAN-based security's primary rule is to confine in-band management and protocol traffic into a controlled environment. This can be achieved with the following tools and best practices:

- Traffic and protocol ACLs or filters.
- QoS marking and prioritization (control protocols are differentiated by means of appropriate class-of-service or DSCP values).
- Selective deactivation of L2 protocols on untrusted ports (for example, disabling DTP on access ports).
- Configuration of inband management port(s) only in dedicated VLAN(s).
- Abstention from using VLAN 1 to carry any data traffic.

Examples of commands:

Catalyst Operating System (CatOS) Software	*Cisco IOS® Software*
`set security acl ip NAME deny tcp any any eq telnet`	`vlan access-map NAME 10`
`set security acl ip NAME permit tcp <host1> any eq ssh`	`match ip address <telnet access list>`
`set security acl ip NAME permit tcp <host2> any eq ssh`	`action drop`
`set security acl ip NAME permit tcp <hostn> any eq ssh`	`vlan access-map NAME 20`
	`match ip address <ssh access list>`
	`action forward`
`set trunk <port range> off`	`switchport mode access (default)`
`set port channel <port range> mode off`	`N/A (default)`
`set cdp disable <port range>`	`no cdp enable`
`set udld disable <port range>`	`udld port disable`
`set interface sc0 <management vlan>`	`interface vlan <management vlan>`
`clear trunk <port range> 1`	`switchport trunk allowed vlan remove 1`

Precautions for the Use of VLAN 1

The reason VLAN 1 became a special VLAN is that L2 devices needed to have a default VLAN to assign to their ports, including their management port(s). In addition to that, many L2 protocols such as CDP, PAgP, and VTP needed to be sent on a specific VLAN on trunk links. For all these purposes VLAN 1 was chosen.

As a consequence, VLAN 1 may sometimes end up unwisely spanning the entire network if not appropriately pruned and, if its diameter is large enough, the risk of instability can increase significantly. Besides the practice of using a potentially omnipresent VLAN for management purposes puts trusted devices to higher risk of security attacks from untrusted devices that by misconfiguration or pure accident gain access to VLAN 1 and try to exploit this unexpected security hole.

To redeem VLAN 1 from its bad reputation, a simple common-sense security principle can be used: as a generic security rule the network administrator *should prune any VLAN, and in particular VLAN 1, from all the ports where that VLAN is not strictly needed*.

Therefore, with regard to VLAN 1, the above rule simply translates into the recommendations to:

- Not use VLAN 1 for inband management traffic and pick a different, specially dedicated VLAN that keeps management traffic separate from user data and protocol traffic.
- Prune VLAN 1 from all the trunks and from all the access ports that don't require it (including not connected and shutdown ports).

  Similarly, the above rule applied to the management VLAN reads:
- Don't configure the management VLAN on any trunk or access port that doesn't require it (including not connected and shutdown ports).
- For foolproof security, when feasible, prefer out-of-band management to inband management. (Refer to [3] for a more detailed description of a out-of-band management infrastructure.)

As a general design rule it is desirable to "prune" unnecessary traffic from particular VLANs. For example, it is often desirable to apply VLAN ACLs and/or IP filters to the traffic carried in the management VLAN to prevent all telnet connections and allow only SSH sessions. Or it may be desirable to apply QoS ACLs to rate limit the maximum amount of ping traffic allowed.

If VLANs other than VLAN 1 or the management VLAN represent a security concern, then automatic or manual pruning should be applied as well. In particular, configuring VTP in transparent or off mode and doing manual pruning of VLANs is commonly considered the most effective method to exert a more strict level of control over a VLAN-based network.

Examples of commands:

CatOS	Cisco IOS Software
`set vtp mode <transparent|off>`	`vtp mode transparent`
`clear trunk <port range> <prunable vlans>`	`switchport trunk allowed vlan remove <prunable vlans>`

"It is an Equal failing to Trust Everybody, and to Trust Nobody" --- English Proverb

After proper handling of VLAN 1 has been decided upon and implemented, the next logical step is to turn one's attention to other equally important best practices commonly used in secure environments. The generic security principle applied here is: *connect untrusted devices to untrusted ports, trusted devices to trusted ports, and disable all the remaining ports*. What this means can be easily expanded into this list of common recommendations:

- If a port is connected to a "foreign" device, don't try to speak any language with it: it could be turned to somebody else's advantage and used against you. So on that port make sure to disable CDP, DTP, PAgP, UDLD, and any other unnecessary protocol, and to enable portfast/BPDU guard on it. After all, why risk a potentially dangerous communication with an untrustworthy neighbor?
- Enable the *rootguard* feature to prevent a directly or indirectly connected STP-capable device to affect the location of the root bridge.
- Configure the VTP domains appropriately or turn off VTP altogether if you want to limit or prevent possible undesirable protocol interactions with regard to network-wide VLAN configuration. This precaution can limit or prevent the risk of an administrator error propagating to the entire network and the risk of a new switch with a higher VTP revision overwriting by accident the entire domain's VLAN configuration.
- By default only those ports which are known to be 'trusted' should be treated as such and all other ports should be configured as 'untrusted'. This prevents attached devices from manipulating QoS values inappropriately.
- Disable unused ports and put them in an unused VLAN. By not granting connectivity or by placing a device into a VLAN not in use, unauthorized access can be thwarted through fundamental physical and logical barriers.

Examples of commands:

CatOS	Cisco IOS Software
`set trunk <port range> off`	`switchport mode access (default)`
`set port channel <port range> mode off`	`N/A (default)`
`set cdp disable <port range>`	`no cdp enable`
`set udld disable <port range>`	`udld port disable`
`set spantree portfast <port range> enable`	`spanning-tree portfast`
`set spantree portfast bpdu-guard <port range> enable`	`spanning-tree portfast bpduguard default`
`set spantree guard root <port range>`	`spanning-tree guard root`
`set vtp mode <transparent|off>`	`vtp mode transparent`
`set port qos <port range> trust untrusted`	`no mls qos trust (default)`
`set port disable <port range>`	`shutdown`

Why Worry About Layer 2 Security in the First Place?

The OSI stack was conceived so that different layers are able to function independently (with only the knowledge of their mutual interfaces). This allows for flexibility in that developments for a given layer of the protocol stack do not impact other layers so long as the standard interface between the layers is maintained.

Unfortunately this also means that if one layer is hacked, communication may be compromised without the other layers being aware of the problem (as shown in *Figure 1*).

Figure 1
OSI Stack Structure

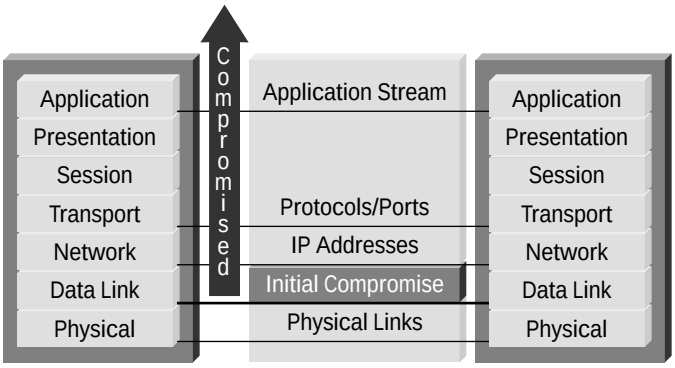

In this architecture, security is only as strong as its weakest link.

The Data Link layer is as vulnerable as any other layer and can be subjected to a variety of attacks which the switch must be configured to protect against.

What Are the Possible Attacks in a VLAN-Based Network?

The majority of attacks at L2 exploit the inability of a device to track the attacker who can therefore perform undetected malicious actions on the forwarding path to alter it and then exploit the change.

These are the most talked-about L2 attacks and incidentally also the ones that @stake documented in its findings [1]:

- MAC Flooding Attack
- 802.1Q and ISL Tagging Attack
- Double-Encapsulated 802.1Q/Nested VLAN Attack
- ARP Attacks
- Private VLAN Attack
- Multicast Brute Force Attack
- Spanning-Tree Attack
- Random Frame Stress Attack

A description of each of these threats follows.

MAC Flooding Attack

This is not properly a network “attack” but more a limitation of the way all switches and bridges work. They possess a finite hardware learning table to store the source addresses of all received packets: when this table becomes full, the traffic that is directed to addresses that cannot be learned anymore will be permanently flooded. Packet flooding however is constrained within the VLAN of origin, therefore no VLAN hopping is permitted (as @ stake’s report shows).

This corner case behavior can be exploited by a malicious user that wants to turn the switch he or she is connected to into a dumb pseudo-hub and sniff all the flooded traffic. Several programs are available to perform this task: for example *macof*, part of the *dsniff* suite [4]. This weakness can then be exploited to perform an actual attack, like the ARP poisoning attack (see *ARP Attacks* for more details on the subject).

On non intelligent switches this problem arises because a sender's L2 identity is not checked, therefore the sender is allowed to impersonate an unlimited number of devices simply by counterfeiting packets.

Cisco's switches support a variety of features whose only goal is to identify and control the identities of connected devices. The security principle on which they are based is very simple: *authentication and accountability are critical for all untrusted devices*.

In particular, *Port Security, 802.1x,* and *Dynamic VLANs* are three features that can be used to constrain the connectivity of a device based on its user's login ID and based on the device's own MAC layer identification.

With Port Security, for instance, preventing any MAC flooding attack becomes as simple as limiting the number of MAC addresses that can be used by a single port: the identification of the traffic of a device is thereby directly tied to its port of origin.

802.1Q and ISL Tagging Attack

Tagging attacks are malicious schemes that allow a user on a VLAN to get unauthorized access to another VLAN. For example, if a switch port were configured as DTP *auto* and were to receive a fake DTP packet, it might become a trunk port and it might start accepting traffic destined for any VLAN. Therefore, a malicious user could start communicating with other VLANs through that compromised port.

Sometimes, even when simply receiving regular packets, a switch port may behave like a full-fledged trunk port (for example, accept packets for VLANs different from the native), even if it is not supposed to. This is commonly referred to as "VLAN leaking" (see [5] for a report on a similar issue).

While the first attack can be prevented very easily by setting DTP to *off* on all non trusted ports (again the principle of trust at work…), the second attack can usually be addressed by following simple configuration guidelines (such as the one suggested in the next section) or with software upgrades. Fortunately, Cisco Catalyst 2950, Catalyst 3550, Catalyst 4000, and Catalyst 6000 series switches don't need any such upgrade, since their software and hardware have been designed to always enforce proper traffic classification and isolation on all their ports (as shown by @ stake in [1]).

Why then is the native VLAN mentioned in the report [5]? The answer is provided in the next section…

Double-Encapsulated 802.1Q/Nested VLAN Attack

While internal to a switch, VLAN numbers and identification are carried in a special extended format that allows the forwarding path to maintain VLAN isolation from end to end without any loss of information. Instead, outside of a switch, the tagging rules are dictated by standards such as ISL or 802.1Q.

ISL is a Cisco proprietary technology and is in a sense a compact form of the extended packet header used inside the device: since every packet always gets a tag, there is no risk of identity loss and therefore of security weaknesses.

On the other hand, the IEEE committee that defined 802.1Q decided that because of backward compatibility it was desirable to support the so-called native VLAN, that is to say, a VLAN that is not associated explicitly to any tag on an 802.1Q link. This VLAN is implicitly used for all the untagged traffic received on an 802.1Q capable port.

This capability is desirable because it allows 802.1Q capable ports to talk to old 802.3 ports directly by sending and receiving untagged traffic. However, in all other cases, it may be very detrimental because packets associated with the native VLAN lose their tags, for example, their identity enforcement, as well as their Class of Service (802.1p bits) when transmitted over an 802.1Q link.

For these sole reasons—loss of means of identification and loss of classification—the use of the native VLAN should be avoided. There is a more subtle reason, though. *Figure 2* shows why.

Figure 2
Double Encapsulation Attack

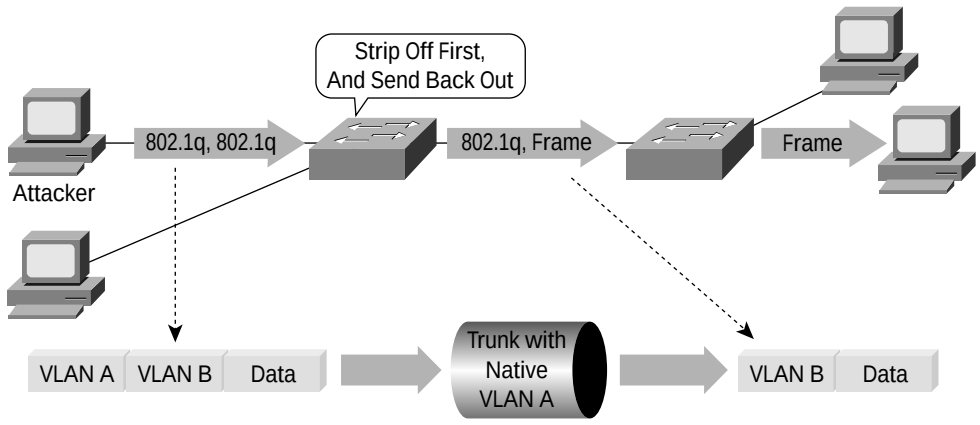

Note: Only Works if Trunk Has the Same Native VLAN as the Attacker.

When double-encapsulated 802.1Q packets are injected into the network from a device whose VLAN happens to be the native VLAN of a trunk, the VLAN identification of those packets cannot be preserved from end to end since the 802.1Q trunk would always modify the packets by stripping their outer tag. After the external tag is removed, the internal tag permanently becomes the packet's only VLAN identifier. Therefore, by double-encapsulating packets with two different tags, traffic can be made to hop across VLANs.

This scenario is to be considered a misconfiguration, since the 802.1Q standard does not necessarily force the users to use the native VLAN in these cases. As a matter of fact, the proper configuration that should always be used is to clear the native VLAN from all 802.1Q trunks (alternatively, setting them to *802.1q-all-tagged* mode achieves the exact same result). In cases where the native VLAN cannot be cleared, *then always pick an unused VLAN as native VLAN of all the trunks; don't use this VLAN for any other purpose*. Protocols like STP, DTP, and UDLD (check out [3]) should be the only rightful users of the native VLAN and their traffic should be completely isolated from any data packets.

ARP Attacks

The ARP protocol [6] is quite an old technology. The ARP RFC is from a time when everyone in a network was supposed to be "friendly" and therefore there was no security built into the ARP function. As a consequence, anyone can claim to be the owner of any IP address they like. To be more precise, anyone can claim that his or her MAC address is associated to any IP address within a specific subnet. This is possible because ARP requests or replies carry the information about the L2 identity (MAC address) and the L3 identity (IP address) of a device and there is no verification mechanism of the correctness of these identities.

Again, this is another case where lack of a precise and reliable means of identification of a device leads to a serious security vulnerability. Also, this is a perfect example of why by compromising a lower level in the OSI stack it's possible to directly affect an upper level without the upper layer being aware of the problem. (ARP is a unique specimen of protocol living and breathing in the L2 world but logically residing at the boundary between the Data Link and the Network layer in the OSI stack.)

The ARP attacks that @stake performed were targeted to fool a switch into forwarding packets to a device in a different VLAN by sending ARP packets containing appropriately forged identities. However, in all Cisco devices VLANs are orthogonal to and therefore independent from MAC addresses: so by changing a device's identity in an ARP packet, it's not possible to affect the way it communicates with other devices *across* VLANs. As a matter of fact, as the report states, any VLAN hopping attempt was thwarted.

On the other hand, *within* the same VLAN, the so-called ARP *poisoning* or *ARP spoofing* attacks [7] are a very effective way to fool end stations or routers into learning counterfeited device identities: this can allow a malicious user to pose as intermediary and perform a *Man-In-the-Middle* (MiM) attack.

In this case, a picture is worth more than a thousand words of explanation (see *Figure 3*).

Figure 3
ARP Poisoning Attack

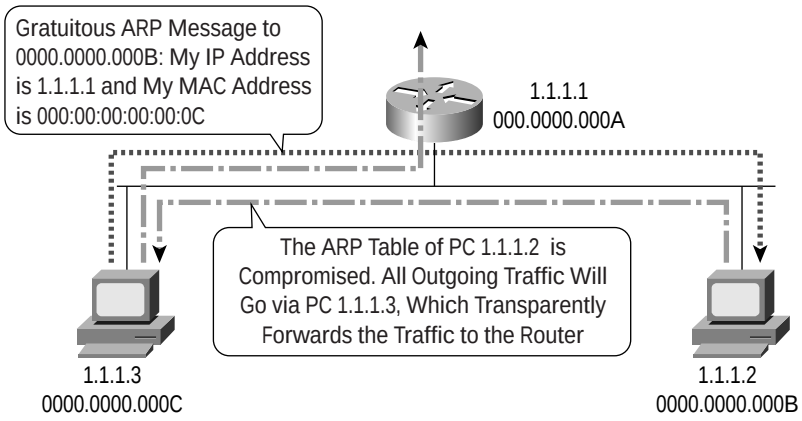

The MiM attack is performed by impersonating another device (for example, the default gateway) in the ARP packets sent to the attacked device: these packets are not verified by the receiver and therefore they "poison" its ARP table with forged information.

This type of attack can be prevented either by blocking the direct communication at L2 between the attacker and the attacked device or by embedding more intelligence into the network so that it can check the forwarded ARP packets for identity correctness. The former countermeasure can be achieved with Cisco Catalyst *Private VLANs* or *Private VLAN Edge* features. The latter can be achieved by using a new feature called ARP Inspection, available first in CatOS 7.5 on the Cisco Catalyst 6500 Supervisor Engine II and a little later also in the Cisco IOS Software for the Cisco Catalyst switches.

Private VLAN Attack

"Private VLAN attack" is actually a misnomer because it corresponds not to a vulnerability but rather to the expected behavior of the feature. Private VLANs is a L2 feature and therefore it is supposed to isolate traffic only at L2. On the other hand, a router is a Layer 3 (L3) device and when it's attached to a Private VLAN promiscuous port it is supposed to forward L3 traffic received on that port to whatever destination it is meant to, even if it's in the same subnet as the source (@stake refers to this behavior as Layer 2 Proxy).

Therefore, it is absolutely normal for two hosts in an Isolated VLAN to fail to communicate with each other through direct L2 communication and instead to succeed to talk to each other by using the router as a packet relay.

Figure 4 depicts the aforementioned behavior.

Figure 4
L2 Proxy

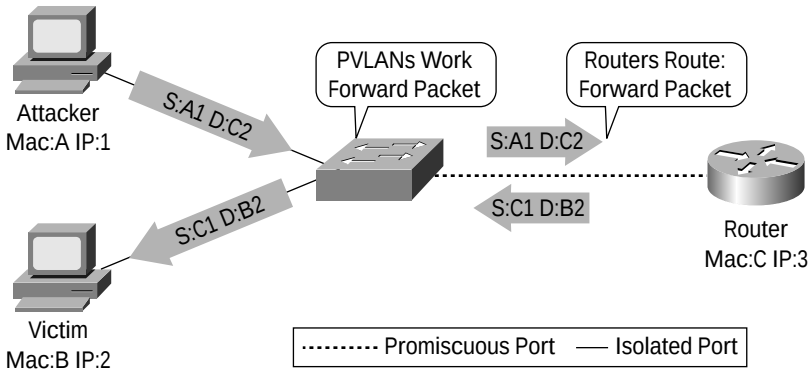

As with regular routed traffic, packets relayed through L2 Proxy can be filtered, if desired, through the configuration of an appropriate ACL on the forwarding device.

Here is a simple example of output Cisco IOS ACL to block the relayed traffic:

```
deny      subnet/mask      subnet/mask
permit      any      subnet/mask
deny      any      any
```

More information on Private VLANs can be found in this paper [8].

Multicast Brute Force Attack

This attack tries to exploit switches' potential vulnerabilities (read: bugs) against a storm of L2 multicast frames. @stake's test was designed to ascertain what happens when a L2 switch receives lots of L2 multicast frames in rapid succession. The correct behavior should be to constrain the traffic to its VLAN of origin, the failure behavior would be to leak frames to other VLANs.

In @stake's results, this type of attack proved ineffective against Cisco Catalyst switches because they correctly contained all the frames within their appropriate broadcast domain (no surprise here: after all, in all Catalyst switches broadcasts are just special cases of multicasts).

Spanning-Tree Attack

Another attack that tries to leverage a possible switch weakness (for example, bug) is the STP attack. All of the Cisco Catalyst switches tested by @stake support this protocol. By default, STP is turned on and every port on the switch both speaks and listens for STP messages. @stake tried to see if Cisco PVST (Per VLAN Spanning Tree) would fail *open across*[1] multiple VLANs under specific conditions. The attack consisted in sniffing for STP frames on the wire to get the ID of the port STP was transmitting on. Next, the attacker would begin sending out STP Configuration/ Topology Change Acknowledgement BPDUs announcing that he was the new root bridge with a much lower priority.

During this procedure broadcast traffic was injected by the testers to discover any possible VLAN leaks, but none were found. This is an indication of the robustness of STP's implementations on Cisco switches.

Random Frame Stress Attack

This last test can have many incarnations but in general it consists in a brute force attack that randomly varies several fields of a packet while keeping only the source and destination addresses constant. After repetitive testing by @stake's engineers, no packets were found to have successfully hopped VLANs.

Private VLANs can be used in this context to better isolate hosts at L2 and shield them from unwanted malicious traffic from untrustworthy devices. Communities of mutually-trusting hosts can be created so as to partition a L2 network into subdomains where only friendly devices are allowed to communicate with each other. For more information on Private VLANs please refer to this paper [8].

Conclusion

The security of VLAN technology has proven to be far more reliable than its detractors had hoped for and only user misconfiguration or improper use of features have been pointed out as ways to undermine its robustness.

The most serious mistake that a user can make is to underestimate the importance of the Data Link layer, and of VLANs in particular, in the sophisticated architecture of switched networks. It should not be forgotten that the OSI stack is only as robust as its weakest link, and that therefore an *equal* amount of attention should be paid to any of its layers so as to make sure that its entire structure is sound.

1. A way in which a device or an algorithm can fail, for example, in a scenario where a device misbehaves and becomes vulnerable and open to attack.

Any good networking design based on Cisco Catalyst switches should incorporate the best practice guidelines described in this paper as an effective way to protect a network's L2 security architecture from dangerous vulnerabilities.

Although some of the security concepts discussed in the previous sections are very generic, this document is solely intended for a network of Cisco Catalyst switches, as other switch vendors' implementations vary greatly and thus some are in fact more susceptible to the various attacks described in this paper.

References

1. Research Report: Secure Use of VLANs: An @stake Security Assessment—August 2002, http://www.cisco.com/warp/public/cc/pd/si/casi/ca6000/tech/stake_wp.pdf
2. SAFE: A Security Blueprint for Enterprise Networks, http://www.cisco.com/go/safe/
3. Best Practices for Catalyst 4500, 5000, and 6500 Series Switch Configuration and Management, http://www.cisco.com/en/US/products/hw/switches/ps663/products_tech_note09186a0080094713.shtml
4. dsniff, by Dug Song, http://monkey.org/~dugsong/dsniff
5. VLAN Security Test Report, July 2000, http://www.sans.org/newlook/resources/IDFAQ/vlan.htm
6. An Ethernet Address Resolution Protocol, RFC 826, http://www.ietf.org/rfc/rfc0826.txt
7. ARP spoofing attack: http://www.sans.org/newlook/resources/IDFAQ/switched_network.htm
8. @stake, http://www.atstake.com/

Acronyms and Definitions

802.1Q	*IEEE specification that defines a standard VLAN tagging scheme.*
BPDU	Bridge Protocol Data Unit *Messages exchanged by switches that run the Spanning Tree Protocol.*
CDP	Cisco Discovery Protocol *Cisco proprietary protocol to discover a network topology made up of compatible devices.*
DTP	Dynamic Trunking Protocol *Cisco proprietary protocol to dynamically negotiate trunking parameters (like status and format).*
IEEE	Institute of Electrical and Electronics Engineers
ISL	Inter-Switch Link *Cisco proprietary VLAN tagging format.*
Native VLAN	*VLAN that is not associated explicitly to any tag on an 802.1Q link.*
OSI	Open Systems Interconnect *Networking Reference Model.*
PAgP	Port Aggregation Protocol *Cisco proprietary protocol to dynamically negotiate channeling parameters (like number of ports).*
STP	Spanning-Tree Protocol *Bridge protocol defined in the IEEE 802.1D standard.*
UDLD	UniDirectional Link Detection *Cisco proprietary protocol to verify the bidirectionality of a physical link.*
VLAN	Virtual Local-Area Network *Virtual broadcast domain comprising one or more switch ports.*
VTP	VLAN Trunking Protocol *Cisco proprietary protocol to distribute VLAN information within a predefined domain.*

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: +65 317 7777
Fax: +65 317 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

MH/LW3985 12/02