

Independent Submission  
Request for Comments: 5517  
Category: Informational  
ISSN: 2070-1721

S. HomChaudhuri  
M. Foschiano  
Cisco Systems  
February 2010

Cisco Systems' Private VLANs:  
Scalable Security in a Multi-Client Environment

Abstract

This document describes a mechanism to achieve device isolation through the application of special Layer 2 forwarding constraints. Such a mechanism allows end devices to share the same IP subnet while being Layer 2 isolated, which in turn allows network designers to employ larger subnets and so reduce the address management overhead.

Some of the numerous deployment scenarios of the aforementioned mechanism (which range from data center designs to Ethernet-to-the-home-basement networks) are mentioned in the following text to exemplify the mechanism's possible usages; however, this document is not intended to cover all such deployment scenarios nor delve into their details.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This is a contribution to the RFC Series, independently of any other RFC stream. The RFC Editor has chosen to publish this document at its discretion and makes no statement about its value for implementation or deployment. Documents approved for publication by the RFC Editor are not a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc5517>.

## Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

## Table of Contents

1. Introduction .....	2
1.1. Security Concerns with Sharing a VLAN .....	3
1.2. The Traditional Solution and Its Related Problems .....	3
2. Private VLANs Architecture .....	4
2.1. VLAN Pairings and Their Port-Related Characteristics .....	7
3. Extending Private VLANs across Switches .....	9
4. A More Flexible IP Addressing Scheme .....	9
5. Routing Considerations .....	10
6. Security Considerations .....	10
7. Acknowledgements .....	11
8. References .....	11
8.1. Normative References .....	11
8.2. Informative References .....	11

## 1. Introduction

In an Ethernet switch, a VLAN is a broadcast domain in which hosts can establish direct communication with one another at Layer 2. If untrusted devices are introduced into a VLAN, security issues may arise because trusted and untrusted devices end up sharing the same broadcast domain.

The traditional solution to this kind of problem is to assign a separate VLAN to each user concerned about Layer 2 security issues. However, the IEEE 802.1Q standard [802.1Q] specifies that the VLAN ID field in an Ethernet frame is 12 bits wide. That allows for a theoretical maximum of 4094 VLANs in an Ethernet network (VLAN numbers 0 and 4095 are reserved). If the network administrator assigns one VLAN per user, then that equates to a maximum of 4094 users that can be supported. The private VLANs technology described in this memo addresses this scalability problem by offering more granular and more flexible Layer 2 segregation, as explained in the following sections.

### 1.1. Security Concerns with Sharing a VLAN

Companies who have Internet presence can either host their servers in their own premises or, alternatively, they can locate their servers at the Internet Service Provider's premises. A typical ISP would have a server farm that offers web-hosting functionality for a number of customers. Co-locating the servers in a server farm offers ease of management but, at the same time, may raise security concerns.

Let us assume that the ISP puts all the servers in one big VLAN. Servers residing in the same VLAN can listen to Layer 2 broadcasts from other servers. Once a server learns the Media Access Control (MAC) address associated to the IP address of another computer in the same VLAN, it can establish direct Layer 2 communication with that device without having to go through a Layer 3 gateway/firewall. If, for example, an attacker gets access to one of the servers, he or she can use that compromised host to launch an attack on other servers in the server farm. To protect themselves from malicious attacks, ISP customers want their machines to be isolated from other machines in the same server farm.

The security concerns become even more apparent in metropolitan area networks. Metropolitan Service Providers may want to provide Layer 2 Ethernet access to homes, rental communities, businesses, etc. In this scenario, the subscriber next door could very well be a malicious network user.

It is therefore very important to offer Layer 2 traffic isolation among customers. Customer A would not want his Layer 2 frames being broadcast to customer B, who happens to be in the same VLAN. Also, customer A would not want customer B to bypass a router or a firewall and establish direct Layer 2 communication with him/her.

### 1.2. The Traditional Solution and Its Related Problems

The traditional solution would be to assign a separate VLAN to each customer. That way, each user would be assured of Layer 2 isolation from devices belonging to other users.

However, with the VLAN-per-customer model, if an ISP wanted to offer web-hosting services to, say, 4000 customers, it would consume 4000 VLANs. Theoretically, the maximum number of VLANs that an 802.1Q-compliant networking device can support is 4094. In reality, many devices support a much smaller number of active VLANs. Even if all devices supported all 4094 VLANs, there would still be a scalability problem when the 4095th customer signed up.

A second problem with assigning a separate VLAN per customer is management of IP addresses. Since each VLAN requires a separate subnet, there can be potential wastage of IP addresses in each subnet. This issue has been described by RFC 3069 [RFC3069] and will not be discussed in detail in this document.

## 2. Private VLANs Architecture

The private VLANs architecture is similar to but more elaborate than the aggregated VLAN model proposed in RFC 3069. The concepts of 'super VLAN' and 'sub VLAN' used in that RFC are functionally similar to the concepts of 'primary VLAN' and 'secondary VLAN' used in this document.

On the other hand, the private VLANs technology differs from the mechanism described in [RFC4562] because instead of using a MAC-address-based 'forced forwarding' scheme it uses a VLAN-based one.

A regular VLAN is a single broadcast domain. The private VLANs technology partitions a larger VLAN broadcast domain into smaller sub-domains. So far, two kinds of special sub-domains specific to the private VLANs technology have been defined: an 'isolated' sub-domain and a 'community' sub-domain. Each sub-domain is defined by assigning a proper designation to a group of switch ports.

Within a private VLAN domain, three separate port designations exist. Each port designation has its own unique set of rules, which regulate a connected endpoint's ability to communicate with other connected endpoints within the same private VLAN domain. The three port designations are promiscuous, isolated, and community.

An endpoint connected to a promiscuous port has the ability to communicate with any endpoint within the private VLAN. Multiple promiscuous ports may be defined within a single private VLAN domain. In most networks, Layer 3 default gateways or network management stations are commonly connected to promiscuous ports.


Isolated ports are typically used for those endpoints that only require access to a limited number of outgoing interfaces on a private-VLAN-enabled device. An endpoint connected to an isolated port will only possess the ability to communicate with those endpoints connected to promiscuous ports. Endpoints connected to adjacent isolated ports cannot communicate with one another. For example, within a web-hosting environment, isolated ports can be used to connect hosts that require access only to default gateways.

A community port is a port that is part of a private VLAN community, which is a grouping of ports connected to devices belonging to the

same entity (for example, a group of hosts of the same ISP customer or a pool of servers in a data center). Within a community, endpoints can communicate with one another and can also communicate with any configured promiscuous port. Endpoints belonging to one community cannot instead communicate with endpoints belonging to a different community or with endpoints connected to isolated ports.

The aforementioned three port designations directly correspond to three different VLAN types (primary, isolated, and community) with well-defined, port-related characteristics, which are described in detail in Section 2.1 below.

Figure 1 below illustrates the private VLAN model from a switch port classification perspective.



- A, B - Isolated devices
- C, D - Community devices
- R - Router (or other L4-L7 device)
- i1, i2 - Isolated switch ports
- c1, c2 - Community switch ports
- p1 - Promiscuous switch port
- t1 - Inter-switch link port (a VLAN-aware port)

Figure 1. Private VLAN classification of switch ports

With reference to Figure 1, each of the port types is described below.

Isolated ports: An isolated port, e.g., i1 or i2, cannot talk to any other port in the private VLAN domain except for promiscuous ports (e.g., p1). If a customer device needs to have access only to a gateway router, then it should be attached to an isolated port.

Community ports: A community port, e.g., c1 or c2, is part of a group of ports. The ports within a community can have Layer 2 communications with one another and can also talk to any promiscuous port. If an ISP customer has, say, 2 devices that he/she wants to be isolated from other customers' devices but to be able to communicate among themselves, then community ports should be used.

Promiscuous ports: As the name suggests, a promiscuous port (p1) can talk to all other types of ports. A promiscuous port can talk to isolated ports as well as community ports and vice versa. Layer 3 gateways, DHCP servers, and other 'trusted' devices that need to communicate with the customer endpoints are typically connected via promiscuous ports.

Please note that isolated, community, and promiscuous ports can either be access ports or hybrid/trunk ports (according to the terminology presented in Annex D of the IEEE 802.1Q specification, up to its 2004 revision).

The table below summarizes the communication privileges between the different private VLAN port types.

	isolated	promiscuous	community1	community2	interswitch link port
isolated	deny	permit	deny	deny	permit
promiscuous	permit	permit	permit	permit	permit
community1	deny	permit	permit	deny	permit
community2	deny	permit	deny	permit	permit
interswitch link port	deny(*)	permit	permit	permit	permit

Table 1

(\*) Please note that this asymmetric behavior is for traffic traversing inter-switch link ports over an isolated VLAN only.

Traffic from an inter-switch link port to an isolated port will be denied if it is in the isolated VLAN. Traffic from an inter-switch link port to an isolated port will be permitted if it is in the primary VLAN (see below for the different VLAN characteristics).

N.B.: An inter-switch link port is simply a regular port that connects two switches (and that happens to carry two or more VLANs).

## 2.1. VLAN Pairings and Their Port-Related Characteristics

In practice, the Layer 2 communication constraints described in the table above can be enforced by creating sub-domains within the same VLAN domain. However, a sub-domain within a VLAN domain cannot be easily implemented with only one VLAN ID. Instead, a mechanism of pairing VLAN IDs can be used to achieve this notion. Specifically, sub-domains can be represented by pairs of VLAN numbers:

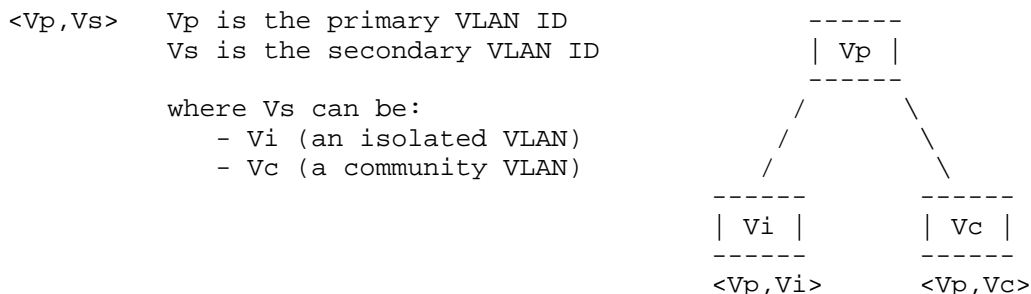


Figure 2. A private VLAN domain can be implemented with one or more VLAN ID pairs.

A private VLAN domain is built with at least one pair of VLAN IDs: one (and only one) primary VLAN ID ( $V_p$ ) plus one or more secondary VLAN IDs ( $V_s$ ). Secondary VLANs can be of two types: isolated VLANs ( $V_i$ ) or community VLANs ( $V_c$ ).

A primary VLAN is the unique and common VLAN identifier of the whole private VLAN domain and of all its VLAN ID pairs.

An isolated VLAN is a secondary VLAN whose distinctive characteristic is that all hosts connected to its ports are isolated at Layer 2. Therefore, its primary quality is that it allows a design based on private VLANs to use a total of only two VLAN identifiers (i.e., a single private VLAN pairing) to provide port isolation and serve any number of end users (vs. a traditional design in which one separate plain VLAN ID would be assigned to each port).

A community VLAN is a secondary VLAN that is associated to a group of ports that connect to a certain "community" of end devices with mutual trust relationships.

While only one isolated VLAN is allowed in a private VLAN domain, there can be multiple distinct community VLANs.

Please note that this VLAN pairing scheme simply requires that all traffic transported within primary and secondary VLANs be tagged according to the IEEE 802.1Q standard (see for example [802.1Q], Section B.1.3), with at most a single standard VLAN tag. No special double-tagging is necessary due to the 1:1 correspondence between a secondary VLAN and its associated primary VLAN.

(Also note that this document makes use of the "traditional" VLAN terminology, whereas the IEEE 802.1ag standard [802.1ag] amends key sections of IEEE 802.1Q-2005 to make the distinction between "VLANs" and "VLAN IDs" so that every "VLAN" can be assigned one or more VLAN IDs, similarly to the pairing scheme described in this document.)

The ports in a private VLAN domain derive their special characteristics (as described in Section 2) from the VLAN pairing(s) they are configured with. In particular, a promiscuous port is a port that can communicate with all other private VLAN port types via the primary VLAN and any associated secondary VLANs, whereas isolated or community ports can communicate over their respective secondary VLANs only.

For example, with reference to Figure 1, a router R connected to the promiscuous port can have Layer 2 communication with a device A connected to an isolated port and also with a device C connected to a community port. Devices C and D can also have Layer 2 communication between themselves since they are part of the same community VLAN. However, devices A and B cannot communicate at Layer 2 due to the special port segregation property of the isolated VLAN. Also, devices A and C cannot communicate at Layer 2 since they belong to different secondary VLANs.

The impact of these enforced forwarding restrictions is two-fold. Firstly, service providers can assign multiple customers to the same isolated VLAN, thereby conserving VLAN IDs. Secondly, end users can be assured that their Layer 2 traffic cannot be sniffed by other end users sharing the same isolated VLAN or connected to a different secondary VLAN.



### 3. Extending Private VLANs across Switches

Some switch vendors have attempted to provide a port isolation feature within a VLAN by implementing special logic at the port level. However, when implemented at the port level, the isolation behavior is restricted to a single switch.

When a VLAN spans multiple switches, there is no standard mechanism to propagate port-level isolation information to other switches and, consequently, the isolation behavior fails in other switches.

In this document, the proposal is to implement the port isolation information implicitly at the VLAN level. A particular VLAN ID can be configured to be the isolated VLAN. All switches in the network would give special "isolated VLAN" treatment to frames tagged with this particular VLAN ID. Thereby, the isolated VLAN behavior can be maintained consistently across all switches in a Layer 2 network.

In general, isolated, community, and primary VLANs can all span multiple switches, just like regular VLANs. Inter-switch link ports need not be aware of the special VLAN type and will carry frames tagged with these VLANs just like they do any other frames.

One of the objectives of the private VLANs architecture is to ensure that traffic from an isolated port in one switch does not reach another isolated or community port in a different switch even after traversing an inter-switch link. By implicitly embedding the isolation information at the VLAN level and by transporting it along with the packet, it is possible to maintain a consistent behavior throughout the network. Therefore, the mechanism discussed in Section 2, which will restrict Layer 2 communication between two isolated ports in the same switch, will also restrict Layer 2 communication between two isolated ports in two different switches.

### 4. A More Flexible IP Addressing Scheme

The common practice of deploying multiple VLANs in a network for security reasons and of allocating a subnet to each VLAN has led to a certain number of inefficiencies in network designs, such as the suboptimal utilization of the IP addressing space (as exemplified in the introduction of RFC 3069 [RFC3069]). Moreover, each subnet requires addresses to be set aside for internetworking purposes (a subnetwork address, a directed broadcast address, default gateway address(es), etc.). So a high number of used VLANs traditionally translates into a significant number of special addresses to be consumed.

On the other hand, in a private VLAN domain, all members can share a common address space that is part of a single subnet associated to the primary VLAN. An end device can be assigned an IP address statically or by using a DHCP server connected to a promiscuous port. Since IP addresses are no longer allocated on a smaller subnet basis but are assigned from a larger address pool shared by all members in the private VLAN domain, address allocation becomes much more efficient: fewer addresses are consumed for internetworking purposes, while most of the address space is allotted to end devices, leaving ample flexibility in the way available addresses are (re-)assigned.

## 5. Routing Considerations

The entire private VLANs architecture confines secondary VLANs within the 2nd layer of the OSI model. With reference to Figure 2, the secondary VLANs are internal to a private VLAN domain. Layer 3 entities are not directly aware of their existence: to them it appears as if all the end devices are part of the primary VLAN.

With reference to Figure 1, the isolation behavior between devices A and B is at the Layer 2 level only. Devices A and B can still communicate at the Layer 3 level via the router R. Since A and B are part of the same subnet, the router assumes that they should be able to talk directly to each other. That however is prevented by the isolated VLAN's specific behavior. So, in order to enable A and B to communicate via the router, a proxy-ARP-like functionality needs to be supported on the router interface.

With regard to the specific version of the IP protocol in use, all routing considerations apply to both IPv4 and IPv6 for the case of unicast traffic. On the other hand, due to their complexity, considerations about multicast bridging and routing within a private VLAN domain transcend the scope of this introductory document, and are therefore omitted.

## 6. Security Considerations

In a heterogeneous Layer 2 network that is built with switches from multiple vendors, the private VLAN feature should be supported and configured on all the switches. If a switch S in that network does not support this feature, then there may be undesired forwarding of packets, including permanent flooding of Layer 2 unicast frames. That is because switch S is not aware of the association between primary and secondary VLANs and consequently cannot apply the segregation rules and constraints characteristic of the private VLANs architecture (an example of one such constraint is explained in [802.1Q], Section B.1.3). This impact is limited to traffic within

the private VLAN domain and will not affect the regular Layer 2 forwarding behavior on other VLANs.

If the private VLAN feature is properly deployed, it can be used at Layer 2 to segregate individual users or groups of users from each other: this segregation allows a network designer to more effectively constrain Layer 2 forwarding so as to, for instance, block or contain unwanted inter-device communication like port scans or Address Resolution Protocol (ARP) poisoning attacks.

## 7. Acknowledgements

Many people have contributed to the private VLANs architecture. We would particularly like to thank, in alphabetical order, Senthil Arunachalam, Jason Chen, Tom Edsall, Michael Fine, Herman Hou, Kannan Kothandaraman, Milind Kulkarni, Heng-Hsin Liao, Tom Nosella, Prasanna Parthasarathy, Ramesh Santhanakrishnan, Mukundan Sudarsan, Charley Wen, and Zhong Xu for their significant contributions.

## 8. References

### 8.1. Normative References

- [802.1Q] Institute of Electrical and Electronics Engineers, "Virtual Bridged Local Area Networks", IEEE Standard 802.1Q, 2005 Edition, May 2006.
- [802.1ag] Institute of Electrical and Electronics Engineers, "Connectivity Fault Management", IEEE Standard 802.1ag, 2007 Edition, December 2007.

### 8.2. Informative References

- [RFC3069] McPherson, D. and B. Dykes, "VLAN Aggregation for Efficient IP Address Allocation", RFC 3069, February 2001.
- [RFC4562] Melsen, T. and S. Blake, "MAC-Forced Forwarding: A Method for Subscriber Separation on an Ethernet Access Network", RFC 4562, June 2006.

Authors' Addresses

Marco Foschiano  
Cisco Systems, Inc.  
Via Torri Bianche 7  
Vimercate, MI, 20059, Italy  
EMail: foschia@cisco.com; mfoschiano@gmail.com

Sanjib HomChaudhuri  
EMail: sanjibhc@gmail.com