

Network Working Group  
Request for Comments: 3069  
Category: Informational

D. McPherson  
Amber Networks, Inc.  
B. Dykes  
Onesecure, Inc.  
February 2001

## VLAN Aggregation for Efficient IP Address Allocation

### Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

### Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

### Abstract

This document introduces the concept of Virtual Local Area Network (VLAN) aggregation as it relates to IPv4 address allocation. A mechanism is described by which hosts that reside in the same physical switched infrastructure, but separate virtual broadcast domains, are addressed from the same IPv4 subnet and share a common default gateway IP address, thereby removing the requirement of a dedicated IP subnet for each virtual Local Area Network (LAN) or Metropolitan Area Network (MAN).


Employing such a mechanism significantly decreases IPv4 address consumption in virtual LANs and MANs. It may also ease administration of IPv4 addresses within the network.

### 1. Introduction

The VLAN [802.1Q] aggregation technique described in this document provides a mechanism by which hosts that reside within the same physical switched infrastructure, but separate virtual broadcast domains, may be addressed from the same IPv4 subnet and may share a common default gateway IPv4 address.

Such a mechanism provides several advantages over traditional IPv4 addressing architectures employed in large switched LANs today. The primary advantage, that of IPv4 address space conservation, can be realized when considering the diagram in Figure 1:

Figure 1:



In the Figure 1 hosts A.1 and A.2 belong to customer A, VLAN A. Hosts B.1 and B.2 belong to customer B, VLAN B. Host C.1 belongs to customer C and resides in it's own virtual LAN, VLAN C.

Traditionally, an IP subnet would be allocated for each customer, based on initial IP requirements for address space utilization, as well as on projections of future utilization. For example, a scheme such as that illustrated in Table 1 may be used.

Table 1:

Customer	IP Subnet	Gateway Address	Usable Hosts	Customer Hosts
A	1.1.1.0/28	1.1.1.1	14	13
B	1.1.1.16/29	1.1.1.17	6	5
C	1.1.1.24/30	1.1.1.25	2	1

Customer A's initial deployment consists of 2 hosts, though they project growth of up to 10 hosts. As a result, they're allocated the IP subnet 1.1.1.0/28 which provides 16 IP addresses. The first IP address, 1.1.1.0, represents the subnetwork number. The last IP address, 1.1.1.15, represents the directed broadcast address. The first usable address of the subnet, 1.1.1.1, is assigned to the router and serves as the default gateway IP address for the subnet. The customer is left 13 IP addresses, even though their requirement was only for 10 IP addresses.

Customer B's initial deployment consists of 2 hosts, though they project growth of up to 5 hosts. As a result, they're allocated the IP subnet 1.1.1.16/29 which provides 8 IP addresses. The first IP address, 1.1.1.16, represents the subnetwork number. The last IP address, 1.1.1.23, represents the directed broadcast address. The first usable address of the subnet, 1.1.1.17, is assigned to the router and serves as the default gateway IP address for the subnet. The customer is left 5 with IP addresses.

Customer C's initial deployment consists of 1 host, and they have no plans of deploying additional hosts. As a result, they're allocated the IP subnet 1.1.1.24/30 which provides 4 IP addresses. The first IP address, 1.1.1.24, represents the subnetwork number. The last IP address, 1.1.1.27, represents the directed broadcast address. The first usable address of the subnet, 1.1.1.25, is assigned to the router and serves as the default gateway IP address for the subnet. The customer is left 1 IP address.

The sum of address requirements for all three customers is 16. The most optimal address allocation scheme here requires 28 IP addresses.

Now, if customer A only grows to use 3 of his available address, the additional IP addresses can't be used for other customers.

Also, assume customer C determines the need to deploy one additional host, and as such, requires one additional IP address. Because all of the addresses within the existing IP subnet 1.1.1.24/30 are used, and the following address space has been allocated to other customers, a new subnet is required. Ideally, the customer would be allocated a /29 and renumber host C.1 into the new subnet. However, the customer is of the opinion that renumbering is not a viable option. As such, another IP subnet is allocated to the customer, this time perhaps a /29, providing two additional addresses for future use.

As you can see, the number of IP addresses consumed by the subnetwork number, directed broadcast address, and a unique gateway address for each subnet is quite significant. Also, the inherent constraints of the addressing architecture significantly reduce flexibility.

## 2. Discussion

If within the switched environment, on the routed side of the network, we introduce the notion of sub-VLANs and super-VLANs, a much more optimal approach to IP addressing can be realized.

Essentially, what occurs is that each sub-VLAN (customer) remains within a separate broadcast domain. One or more sub-VLANs belong to a super-VLAN, and utilize the default gateway IP address of the super-VLAN. Hosts within the sub-VLANs are numbered out of IP subnets associated with the super-VLAN, and their IP subnet masking information reflects that of the super-VLAN subnet.

If desired, the super-VLAN router performs functions similar to Proxy ARP to enable communication between hosts that are members of different sub-VLANs.

This model results in a much more efficient address allocation architecture. It also provides network operators with a mechanism to provide standard default gateway address assignments.

Let's again consider Figure 1, now utilizing the super-VLAN sub-VLAN model. Table 2 provides the new addressing model.

Table 2:

Customer	IP Subnet	Gateway Address	Usable Hosts	Customer Hosts
=====	=====	=====	=====	=====
A	1.1.1.0/24	1.1.1.1	10	.2-.11
B	1.1.1.0/24	1.1.1.1	5	.12-.16
C	1.1.1.0/24	1.1.1.1	1	.17

Customer A's initial deployment consists of 2 hosts, though they project growth of up to 10 hosts. As a result, they're allocated the IP address range 1.1.1.2 - 1.1.1.11. The gateway address for the customer is 1.1.1.1, the subnet is 1.1.1.0/24.

Customer B's initial deployment consists of 2 hosts, though they project growth of up to 5 hosts. As a result, they're allocated the IP address range 1.1.1.12 - 1.1.1.16. The gateway address for the customer is 1.1.1.1, the subnet is 1.1.1.0/24.

Customer C's initial deployment consists of 1 host, and they have no plans of deploying additional hosts. As a result, they're allocated the IP address 1.1.1.17. The gateway address for the customer is 1.1.1.1, the subnet is 1.1.1.0/24.

The sum of address requirements for all three customers is 16. As a result, only 16 addresses are allocated within the subnet. These 16 addresses, combined with the global default gateway address of 1.1.1.1, as well as the subnetwork number of 1.1.1.0 and directed broadcast of 1.1.1.255, result in a total of 19 addresses used. This leaves 236 additional usable hosts address with the IP subnet.

Now, if customer A only grows to use 3 of his available addresses, the additional IP addresses can be used for other customers.

Also, assume customer C determines the need to deploy one additional host, and as such, requires one additional IP address. The customer is simply allocated the next available IP address within the subnet, their default gateway remains the same.

The benefits of such a model are obvious, especially when employed in large LANs or MANs.

### 3. Use of Directed Broadcasts

This specification provides no support for directed broadcasts. Specifically, the <net, subnet, -1> directed broadcast address can only apply to one of the Layer 2 broadcast domains.

Though use of directed broadcast is frowned upon in the Internet today, there remain a number of applications, primarily in the enterprise arena, that continue to use them. As such, care should be taken to understand the implications of using these applications in conjunction with the addressing model outlined in this specification.

### 4. Multicast Considerations

It is assumed that the Layer 2 multicast domain will be the same as the Layer 2 broadcast domain (i.e., VLAN). As such, this means that for an IP multicast packet to reach all potential receivers in the IP subnet the multicast router(s) attached to the IP subnet need to employ something akin to IP host routes for the sender in order for the Reverse Path Forwarding check to work.

### 5. Deployment Considerations

Extreme Networks has a working implementation of this model that has been deployed in service provider data center environments for over a year now. Other vendors are rumored to be developing similar functionality.

### 6. Security Considerations

One obvious issue that does arise with this model is the vulnerabilities created by permitting arbitrary allocation of addresses across disparate broadcast domains. It is advised that address space ranges be made sticky. That is, when an address or range of addresses is allocated to a given sub-VLAN, reception of IP

or ARP packets on a sub-VLAN with a source IP address that isn't allocated to the sub-VLAN should be discarded, and perhaps trigger a logging message or other administrative event.

Implementation details are intentionally omitted as all functions in this document should remain local to the super-VLAN router. As such, no interoperability issues with existing protocols should result.

## 7. Acknowledgements

Thanks to Mike Hollyman and Erik Nordmark for their feedback.

## 8. References

[802.1Q] IEEE 802.1Q, "Virtual LANs".

## 9. Authors' Addresses

Danny McPherson  
Amber Networks, Inc.  
48664 Milmont Drive  
Fremont, CA 94538

EEmail: danny@ambernetworks.com

Barry Dykes  
OneSecure, Inc.  
2000 S. Colorado Blvd Suite 2-1100  
Denver, CO. 80222

EEmail: bdykes@onesecure.com

## 10. Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

## Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.