WHITE PAPER

VLAN DESIGN FOR IPTV/MULTIPLAY NETWORKS

Understanding the Alternatives

Table of Contents

Executive Summary 1
Introduction 1
VLAN Overview 1
VLAN Design Alternatives 2
Service VLAN (S-VLAN) Model 2
Strengths 2
Weaknesses 2
Customer VLANs 3
Strengths 3
Weaknesses 4
C-VLAN: Adding an Intermediate Aggregation Switch 4
Hybrid: Customer VLAN (C-VLAN) with Multicast VLAN (M-VLAN) 5
Strengths 5
Weaknesses 5
Last Mile Connection 5
Juniper Networks VLAN Support for IPTV/Multiplay 6
Conclusion 6
About Juniper Networks 6

Table of Figures

Figure 1: Ethernet frame with VLAN tag 1
Figure 3: Simple bandwidth carve-out 2
Figure 2: Service VLAN model 2
Figure 4: Customer VLAN model 3
Figure 5: Dynamic bandwidth allocation using customer VLAN 3
Figure 6: IPTV traffic in a C-VLAN network 4
Figure 7: C-VLAN Ethernet frame with intermediate switch 4
Figure 8: Hybrid (C-VLAN with M-VLAN) model 5
Figure 9: Last mile connection 5

Executive Summary

This document discusses the two major VLAN network design alternatives for supporting multiplay (IPTV and unicast) services across a broadband network. It is intended for network planners and operations personnel looking to offer multiple services across a wireline broadband access network such as xDSL or a passive optical network (PON).

Introduction

VLANs limit the broadcast domain, ensuring that subscribers cannot see each other's information and also reducing network traffic. They are a fundamental part of how broadband access networks deliver information to subscribers. There are two ways that VLANs are used in this role. A service VLAN (S-VLAN) delivers a single service to all subscribers, while a customer VLAN (C-VLAN) delivers multiple services to a single subscriber. These are often combined by using a single S-VLAN to carry multicast traffic to all subscribers, but using dedicated C-VLANs to carry all unicast traffic to each subscriber. Each of these scenarios differs from the others in manageability, bandwidth requirements and deployability.

VLAN Overview

VLANs allow networks to logically segment into multiple communities by "tagging" each packet with a field which identifies packet ownership. VLAN usage is standardized as IEEE 802.1Q, and tags are sometimes called "VLAN tags" or "Q-tags." Figure 1 depicts a standard Ethernet frame with the optional VLAN tag included in the packet. For this discussion, the critical field is the 12-bit VLAN identifier (VLAN id), which identifies the user community to which this packet belongs. A single VLAN tag supports up to 4,095 (212, with VLAN id = 0 reserved) user communities.

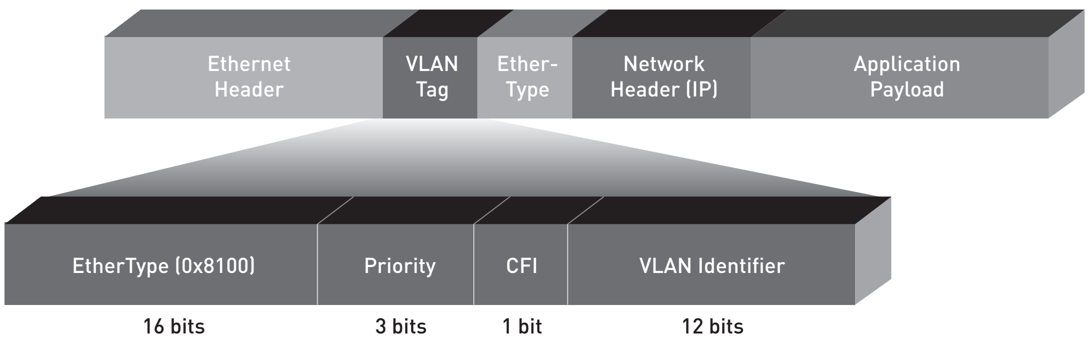

Figure 1: Ethernet frame with VLAN tag

An important attribute of each VLAN is that it is a unique Ethernet broadcast domain. This means that traffic tagged as belonging to one VLAN group will be seen only by authorized subscribers who are also members of this group. VLANs were originally designed for use within a single organization where 4,095 VLAN identifiers were considered sufficient. Service Providers also use VLANs to differentiate between customers.

VLAN Design Alternatives

In broadband networks, there are two fundamental VLAN design alternatives:

- Customer VLAN: Also called the 1:1 model, in this model there is a dedicated VLAN for each subscriber.
- Service VLAN: In this model there is a dedicated VLAN for each service. This is also called the N:1 model since multiple subscribers share each VLAN.

Service VLAN (S-VLAN) Model

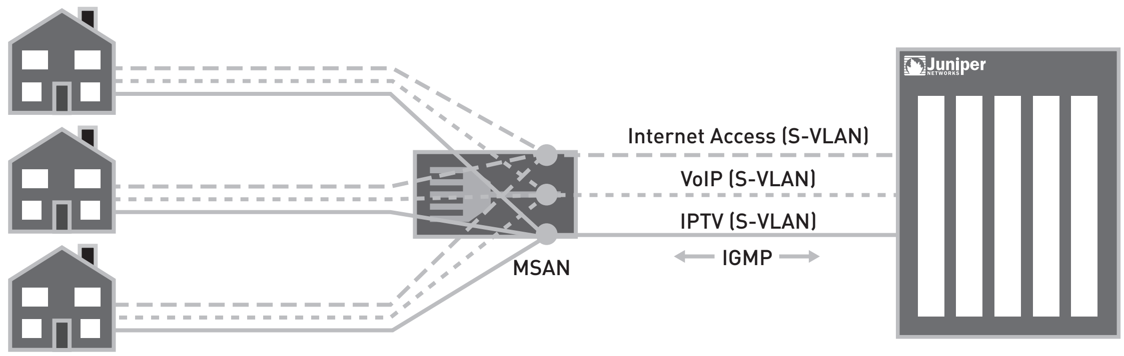

Figure 2: Service VLAN model

In the service VLAN model, there is a separate VLAN for each service such as Internet access, Voice over IP (VoIP), IP-based broadcast television (IPTV) and Video on Demand (VoD). Internet Group Management Protocol (IGMP) packets always travel on the same S-VLAN as the associated IPTV. Figure 2 shows a simple example with three service VLANs.

The big challenge to service VLANs is that no network element can look at all traffic destined for each specific subscriber to determine whether there is sufficient bandwidth. Instead, a fixed amount of bandwidth must be “carved-out” for each service. For example, the bandwidth carve-out depicted in Figure 3 allows the subscriber to have:

- Up to 8 Mbps of video that supports up to 4 SDTVs (@ 2 Mbps, or 1 HDTV stream (@ 6 Mbps) plus one SDTV
- Up to 2 Mbps for data

However, it does not permit this mix to be altered, such as by allowing a telecommuting power user to have a single SDTV connection and utilize the remaining bandwidth for data downloads. Conversely, it also does not allow the data bandwidth to be used to deliver video to additional screens, an important requirement as HD penetration grows.

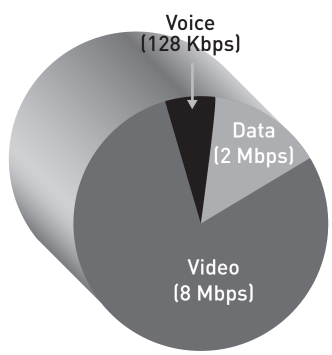

Figure 3: Simple bandwidth carve-out

Service VLANs are typically used by operators looking to offer a few basic services with the simplest roll-out and minimal investment.

Strengths

- A new service can be added by creating a new service VLAN. This is the simplest to provision if the network cannot dynamically create subscriber instances (C-VLANs).
- Some multiservice access nodes (MSANs) only support this model because they can carry a limited number of VLANs.
- Every residential gateway—the device within the home which connects to the DSL line—has the same configuration.

Weaknesses

- In the home, each service is mapped to a specific RG port. This makes it difficult to separate “services” from “devices,” for instance by offering IPTV service on a PC.
- Call admission control cannot be implemented because of the intensive processing power that would be required by the MSAN to inspect each packet in each service VLAN to determine the total traffic flowing to each subscriber.

Customer VLANs

The customer VLAN (C-VLAN) model as depicted in Figure 4 uses a separate dedicated VLAN for each subscriber. This C-VLAN carries all traffic between the multiservice access nodes (MSANs), such as an optical line termination (OLT) or digital subscriber line access multiplexor (DSLAM) and the broadband services router (BSR). This model mirrors the deployed edge architecture used in many carrier environments for dial, private line, Metro Ethernet and Frame/ATM aggregation, and is deployed by many of the world's largest IPTV providers.

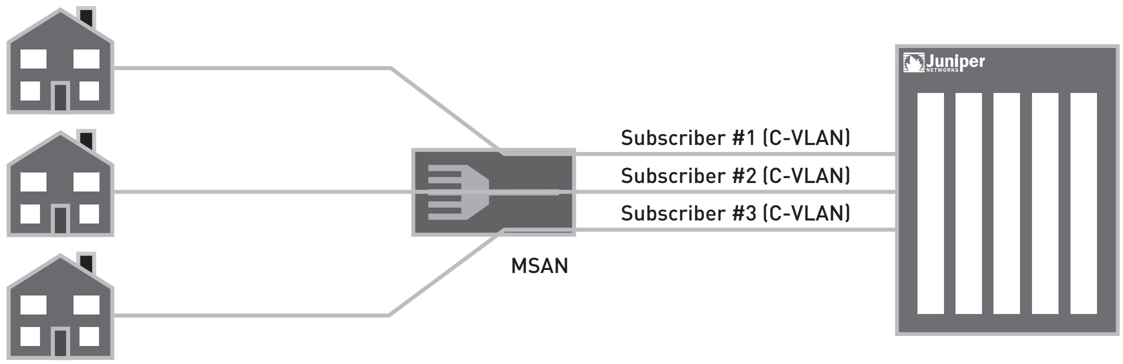

Figure 4: Customer VLAN model

In a "pure" C-VLAN model, the edge router performs channel replication and sends the broadcast television traffic to the subscriber via a unicast session. IPTV traffic is forwarded to each subscriber across the C-VLAN. A pure C-VLAN model works well if most television traffic is expected to be unicast to each subscriber (VoD or internet-based video downloads), or if there is sufficient bandwidth available to send a unique video stream to each subscriber.

Strengths

- The key strength of the C-VLAN model is that it enables the edge router to effectively manage the bandwidth for each subscriber and implement call admission control. Since there is a single, intelligent device which sees all traffic, it can determine whether there is sufficient bandwidth available to allow the subscriber to access another application.
- The MSAN becomes a simple cross-connect device, allowing a common operational model regardless of the selected MSAN, and allowing the operator to purchase the lowest priced MSAN.
- Mapping each subscriber to a single VLAN and eliminating multicast in the access network simplifies problem determination.

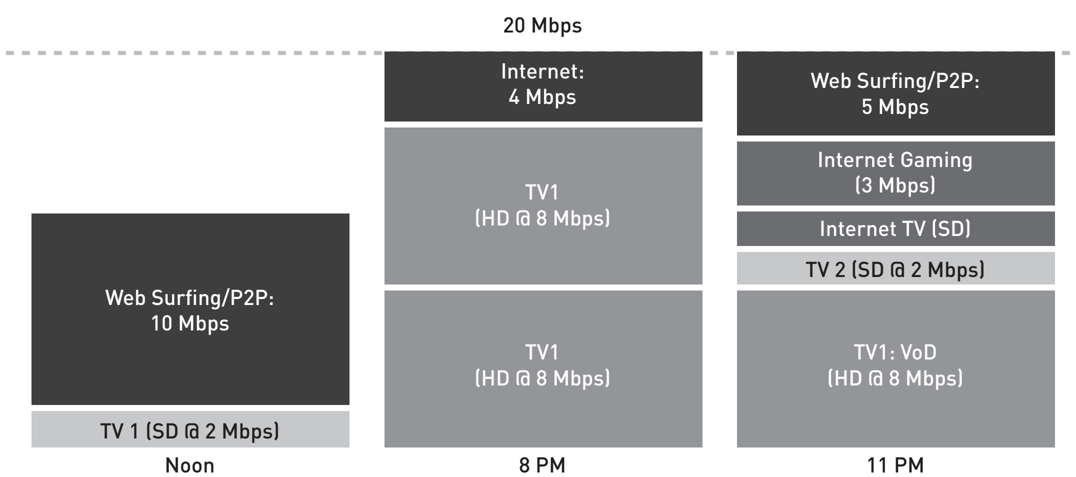

Figure 5: Dynamic bandwidth allocation using customer VLAN

Weaknesses

- This "pure C-VLAN" model increases the bandwidth required between the BSR and the MSAN, since a separate copy of each channel being viewed is be sent to the MSAN. As illustrated in Figure 6, the edge router would need to send a separate IPTV stream to each subscriber, even those watching the same channel. The MSAN cannot replicate channels since it does not realize that the same channel is coming in over multiple virtual connections.
- Some MSANs do not support the larger number of VLANs dictated by this model.
- In the simplest configurations, the Ethernet frame is passed transparently through the MSAN. As a result, each RG must be configured to support different VLAN-id. (Some MSANs can delete or replace the VLAN tag before forwarding the packet to the RG, eliminating this issue).

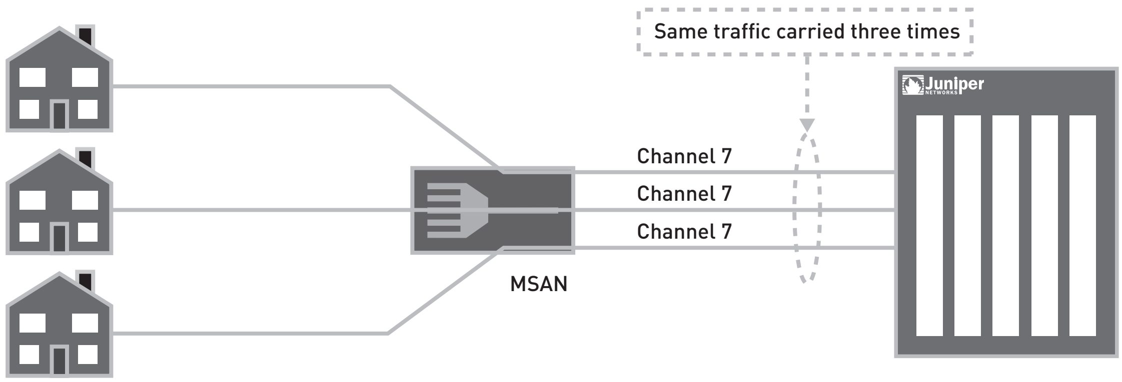

Figure 6: IPTV traffic in a C-VLAN network

C-VLAN: Adding an Intermediate Aggregation Switch

MSANs are often connected to the edge router via an intermediate Ethernet aggregation switch. One way to support this is by assigning a unique VLAN id to each subscriber connected to the same edge router port. However, since a single VLAN tag supports only 4,095 subscribers, it is possible to deplete this pool.

As depicted in Figure 7, the solution is to use multiple VLAN tags—an "outer" tag and and "inner" tag. The outer VLAN tag identifies the MSAN and is used by the switch to forward traffic to it. In addition, the switch removes this tag. The inner tag identifies the subscriber connected to this MSAN and is used to identify the individual subscriber. This technique is defined in the IEEE 802.1ad supplement to the IEEE 802.1Q VLAN Bridging standard, and is commonly called "Q-in-Q" or "stacked VLANs".

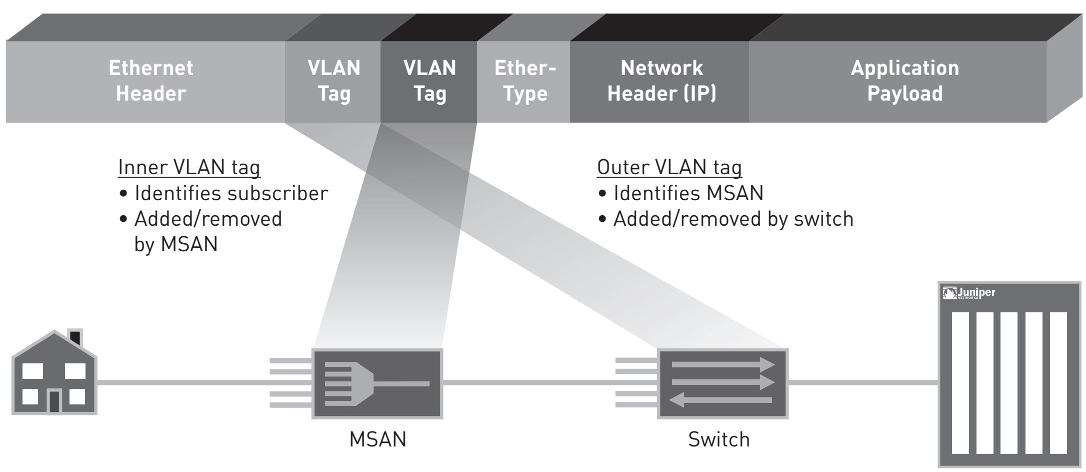

Figure 7: C-VLAN Ethernet frame with intermediate switch

Hybrid: Customer VLAN (C-VLAN) with Multicast VLAN (M-VLAN)

The hybrid model leverages the strengths of both architectures, creating a single policy enforcement point while providing efficient multicast delivery. The hybrid model as depicted in Figure 8 leverages multiple VLANs as follows:

- A subscriber-dedicated C-VLAN carries unicast traffic such as Internet Access and Voice over IP between the access node and the BSR.
- A service VLAN carries broadcast television traffic to each MSAN. Because the S-VLAN carries multicast IPTV traffic exclusively, it is often referred to as a Multicast VLAN (M-VLAN).

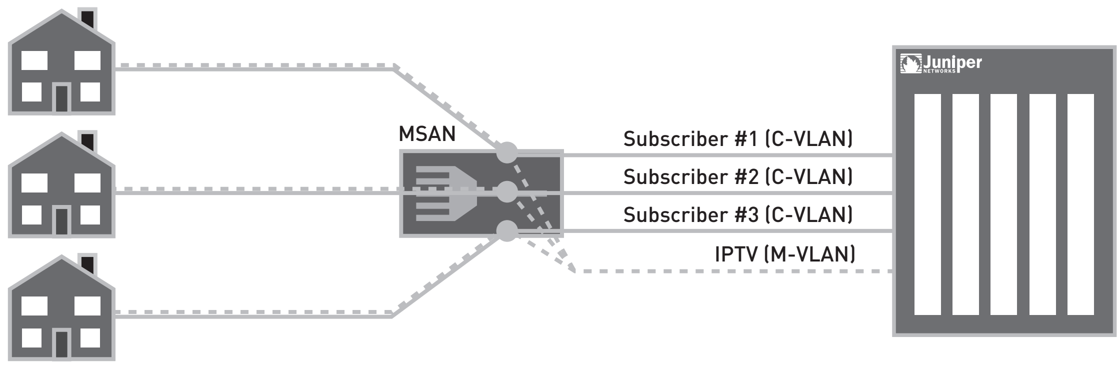

Figure 8: Hybrid (C-VLAN with M-VLAN) model

In this model, there is one "customer VLAN" per subscriber which carries all unicast traffic (Internet access, VoIP, VoD). In addition, one VLAN carries multicast traffic to all subscribers. Therefore, if there are N subscribers, the network supports N+1 VLANs.

At the MSAN, the separate VLANs are often merged onto a single VLAN. IGMP and VoD traffic can flow on either the C-VLAN or the M-VLAN. IGMP Forking can also be used to forward IGMP requests on both VLANs or ATM VCs.

Strengths

This model provides the simple subscriber VLAN separation to simplify problem determination. In addition, the BSR can track the amount of bandwidth to each subscriber required to deliver IPTV traffic. This allows the BSR to adjust bandwidth available to each application in each C-VLAN and implement call admission control (CAC) as required.

Weaknesses

The MSAN must be able to inspect incoming (upstream) traffic and add the appropriate VLAN tag.

Last Mile Connection

Figure 9 depicts the alternatives when using a C-VLAN model. Ideally, all traffic from the MSAN to the RG typically includes a VLAN id of 0 (top diagram). The VLAN tag is included to allow the RG to specify upstream traffic priority by setting the 802.1p priority bits. Having the same VLAN id for all subscribers allows a single RG configuration to be used. Some MSANs cannot replace the VLAN tags, so the existing VLAN tags are carried to the RG (bottom).

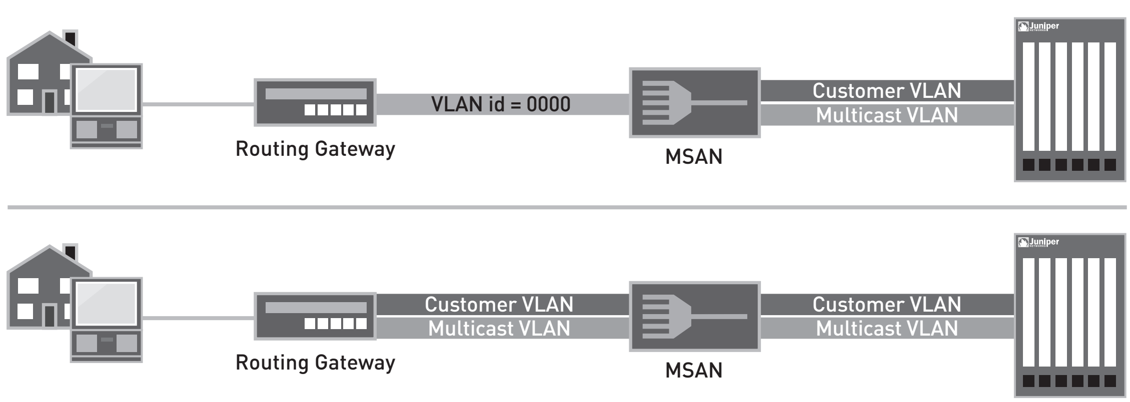

Figure 9: Last mile connection

Juniper Networks VLAN Support for IPTV/Multiplay

Juniper Networks® broadband services routers support the complete range of capabilities supporting IPTV over broadband networks. They support all of the models discussed, including stacked VLANs. Juniper's solution also provides call admission control for unicast and multicast traffic, extensive quality of service (QoS) support and advanced hierarchical queuing to ensure that each application is successfully delivered to each subscriber.

Conclusion

All three topologies described in this paper are deployed today. For a standard retail broadband network with a large amount of broadcast television, the S-VLAN model works well and is often used. In addition, the service VLAN model is typically more appealing tor service providers looking to build a multi-edge network or provide wholesaler-based IPV service, since VLANs carrying different services can easily be forwarded to different endpoints. For networks with a heavy focus on offering diverse services or heavy VoD penetration, the C-VLAN model is preferred. A separate M-VLAN is deployed if dictated by bandwidth requirements.

About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at **www.juniper.net.**

Corporate And Sales Headquarters
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER
(888.586.4737)
or 408.745.2000
Fax: 408.745.2100

APAC Headquarters
Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin,
Ireland
Phone: 35.31.8903.600
Fax: 35.31.8903.601

2000186-001-EN Feb 2009

Printed on recycled paper.