A Survey of Virtual LAN Usage in Campus Networks

Minlan Yu and Jennifer Rexford, Princeton University
Xin Sun and Sanjay Rao, Purdue University
Nick Feamster, Georgia Institute of Technology

ABSTRACT

VLANs are widely used in today's enterprise networks to improve Ethernet scalability and support network policies. However, manuals and textbooks offer very little information about how VLANs are actually used in practice. Through discussions with network administrators and analysis of configuration data, we describe how three university campuses and one academic department use VLANs to achieve a variety of goals. We argue that VLANs are ill-suited to some of these goals (e.g., VLANs are often used to realize access control policies, but constrain the types of policies that can be expressed). Furthermore, the use of VLANs leads to significant complexity in the configuration of network devices.

INTRODUCTION

Enterprise networks, which connect the computers within a college campus or corporate location, differ markedly from backbone networks. These networks have distinctive topologies, protocols, policies, and configuration practices. Yet, the unique challenges in enterprise networks are not well understood outside of the operator community. One prominent example is virtual LANs (VLANs) — a widely- used technology that is barely discussed in networking textbooks.

VLANs were initially intended to allow network administrators to connect a group of hosts in the same broadcast domain, independent of their physical location. However, today's enterprise administrators use VLANs for a variety of other purposes, most notably for better scalability and flexible specification of policies. However, enterprise administrators have seen many problems of VLANs because VLANs are used for other functions they were not designed for. Understandably, VLANs are at best an incomplete solution for some of these problems. As a result, managing VLANs is one of the most challenging tasks they face.

In this article, we study four networks — three university campuses and one academic department — to better understand how VLANs are used in practice. Through discussions with network administrators, and targeted analysis of router configuration data, we have obtained deeper insights into how the administrators use VLANs to achieve a variety of design goals, and the difficulties they encounter in the process. We show that VLANs are *not* well-suited for many of the tasks that they support today, and argue that future enterprise network architectures should decouple policy specification from scalability concerns with layer-2 protocols, topology, and addressing.

After a brief survey of VLAN technology, we describe how the four networks use VLANs to support resource isolation, access control, decentralized management, and host mobility. However, VLANs were not designed with these goals in mind — network administrators use VLANs for the lack of a better alternative. We argue that VLANs are too crude a mechanism for specifying policies, due to *scalability* constraints (on the number and size of VLANs) and the *coarse-grained* ways of assigning traffic to different VLANs. Further, VLAN configuration is far too complicated, due to the tight coupling with spanning-tree construction, failure recovery, host address assignment, and IP routing, as discussed. We conclude the article.

VIRTUAL LOCAL AREA NETWORKS

An enterprise network consists of islands of Ethernet switches connected both to each other and to the rest of the Internet by IP routers, as shown in Fig. 1. We describe how administrators group related hosts into VLANs, and how the switches and routers forward traffic between hosts.

CONVENTIONAL LOCAL AREA NETWORKS

In a traditional local area network (LAN), hosts are connected by a network of hubs and switches. The switches cooperate to construct a *spanning tree* for delivering traffic. Each switch forwards Ethernet frames based on its destination MAC address. If the switch contains no forwarding-table entry for the frame's destination MAC address, the switch *floods* each frame over the entire spanning tree. A switch *learns* how to reach a MAC address by remembering the

0163-6804/11/$25.00 © 2011 IEEE

incoming link for frames sent by that MAC address and creating a mapping between the MAC address and that port.

To connect to the rest of the enterprise network (and the rest of the Internet), the island of Ethernet switches connects to IP routers that forward traffic to and from remote hosts. Each host interface in the LAN has an IP address from a common IP prefix (or set of prefixes). Traffic sent to an IP address in the same subnet stays within the LAN; the sending host uses the Address Resolution Protocol (ARP) to determine the MAC address associated with the destination IP address. For traffic destined to remote IP addresses, the host forwards the packets to the gateway router, which forwards packets further toward their destinations.

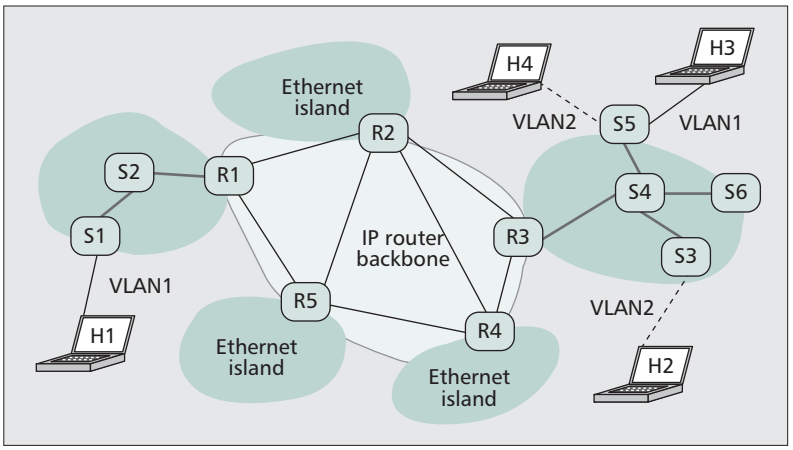

Figure 1. *Enterprise network with Ethernet islands interconnected by IP routers.*

Communication Within A VLAN

Administrators use VLANs to construct network segments that behave logically like a conventional LAN but are independent of the physical locations of the hosts; for example, hosts H1 and H3 in Fig. 1 both belong to VLAN1. As in a conventional physical LAN, the switches in a VLAN construct a spanning tree, and use flooding and learning to forward traffic between hosts. For example, the switches S3, S4, and S5 form a spanning tree for VLAN2.

Communication between hosts in the same VLAN stays within the VLAN, with the switches forwarding Ethernet frames along the spanning tree to the destination MAC address. For example, hosts H2 and H4 communicate over the 2 spanning tree in VLAN2 based on their MAC addresses. Similarly, hosts H1 and H3 communicate over the spanning tree in VLAN1, where some of the IP routers (e.g., R1, R2, and R2) may also act as switches in the spanning tree; alternatively, a tunnel between R1 and R2 could participate in VLAN1 so the links in the IP backbone do not need to participate in the VLANs.

Communication Bbetween VLANs

Each host has an IP address from an IP prefix (or prefixes) associated with its VLAN; IP routers forward packets based on these prefixes, over paths computed in the routing protocol (e.g., Open Shortest Path First [OSPF] or Routing Information Protocol [RIP]). Hence, traffic between hosts in different VLANs must traverse an intermediate IP router. For example, traffic between hosts H3 and H4 would traverse router R2, *even though the two hosts connect to the same switch*. For example, when sending traffic to H4, host H3 forwards the packets to its gateway router R2, since the destination IP address belongs to a different prefix. R2 would then look up the destination IP address to forward the packet to H4 in VLAN2. If H4 sends an IP packet to H1, then H4's router R3 forwards the packet based on the IP routing protocol toward the router announcing H1's IP prefix, and that router would then forward the packet over the spanning tree for VLAN1.

Configuring VLAN Ports

Supporting VLANs requires a way to associate switch ports with one or more VLANs. Administrators configure each port as either an *access port*, which is connected to a host; or a *trunk port*, which is connected to another switch. An access port typically transports traffic for a single VLAN; the VLAN associated with a port may be either statically configured or dynamically assigned when the host connects, based on the host's MAC address (e.g., using VLAN Management Policy Server VMPS [1]). In either case, the access port can tag incoming frames with the 12-bit VLAN identifier and removes the tag from outgoing frames, obviating the need for the hosts to support VLANs.

In contrast, a trunk port may carry traffic for multiple VLANs; for example, switch S4's port connecting to S5 must forward traffic for both VLAN1 and VLAN2 (and participate in each VLAN's spanning tree protocol), but the trunk port to S3 does not. The administrators either manually configure each trunk port with a list of VLAN identifiers, or run a protocol like VLAN Trunking Protocol (VTP) [2] or Multiple VLAN Registration Protocol (MVRP) [3] to automatically determine which VLANs a trunk link should handle. Configuring a VLAN also requires configuring the gateway router to announce the associated IP prefixes into the routing protocol; each host interface must be assigned an IP address from the prefix associated with its VLAN.

VLAN Usage in Campus Networks

Our campus network administrators use VLANs to achieve four main policy objectives — limiting the scope of broadcast traffic, simplifying access control policies, supporting decentralized network management, and enabling seamless host mobility for wireless users. The four networks include two large universities (campuses 1 and 2) and a department network (campus 3) within another university-wide network (campus 4). All four networks primarily run IPv4, with relatively limited experimental deployment of IPv6.

Scoping Broadcast Traffic

VLANs enable administrators to limit the scope of broadcast traffic and network-wide flooding, to reduce network overhead and enhance both privacy and security.

VLANs provide an effective way to enforce access control by directing inter-VLAN traffic through routers. In addition, by allowing administrators to assign related hosts to IP addresses in the same subnet, VLANs simplify access control configuration by making packet-classification rules more concise.

Limiting the Broadcast/Flooding Overhead — End hosts broadcast Dynamic Host Configuration Protocol (DHCP) traffic when joining the LAN, and routinely broadcast Address Resolution Protocol (ARP) requests to learn the medium access control (MAC) addresses of other hosts in the same IP subnet. For example, campus 2 has one IP subnet with up to 4000 hosts with around 300 packets/s of broadcast traffic; this broadcast traffic is dominated by ARP, iTunes broadcast messages, and NetBios. It not only consumes network bandwidth, but also consumes bandwidth and energy resources on the *end hosts* (particularly for mobile devices). Switches also flood packets to a destination MAC address they have not yet learned how to reach. This consumes bandwidth resources, especially if the switches' forwarding tables are not large enough to store an entry for each MAC address on the LAN. Administrators often divide large networks into multiple VLANs to limit the scope of broadcast messages and flooding traffic. For example, campuses 1 and 4 assign each building a different IP subnet, each associated with its own VLAN. The resulting broadcast domains are small enough to limit the overhead on the switches and the end hosts.

Protecting Security and Privacy — Broadcast and flooding traffic also raise security and privacy concerns. Sending excessive broadcast traffic is an effective denial-of-service attack on the network. In addition, a malicious host can intentionally overload switch forwarding tables (e.g., by spoofing many source MAC addresses), forcing switches to flood legitimate traffic that can be easily monitored by the attacking host. ARP is also vulnerable to man-in-the-middle attacks, where a malicious host sends unsolicited ARP responses to impersonate another host on the LAN, thereby intercepting all traffic sent to the victim. Network administrators can reduce these risks by constraining which users can belong to the same VLAN. For example, campus 3 has separate subnets for faculty, graduate students, and undergraduate students, and assigns each subnet to one VLAN based on the registered MAC addresses of the user machines. This ensures that students cannot intercept faculty traffic (e.g., a midterm exam en route to the printer), and that research experiments on the graduate-student VLAN do not inadvertently overload the faculty VLAN.

Simplifying Access Control Policies

VLANs provide an effective way to enforce access control by directing inter-VLAN traffic through routers. In addition, by allowing administrators to assign related hosts to IP addresses in the same subnet, VLANs simplify access control configuration by making packet classification rules more concise.

Imposing Access Control Policies — VLANs provide a way to restrict communication between hosts. In Fig. 1, router 3 (R3) can apply access control lists (ACLs) to limit the traffic between hosts H3 and H4 that belong to different VLANs. Along the same lines, administrators do not place hosts in the same VLAN *unless* they are allowed to communicate freely. Campus 3, for example, places all infrastructure services — such as e-mail and DHCP servers — on a single VLAN since these managed services all trust each other. As another example, campus 1 has several "private" VLANs that have *no* IP router connecting them to the rest of the IP network; for example, the automatic teller machines (ATMs) belong to a private VLAN to protect them from attacks by other hosts.

Concise Access Control Lists — Routers and firewalls apply ACLs based on the five-tuple of the source and destination IP addresses, the source and destination TCP/UDP port numbers, and the protocol. Wildcards enable shorter lists of rules for permitting and denying traffic, which simplifies ACL configuration and also makes efficient use of the limited high-speed memory (e.g., TCAMs) for applying the rules. VLANs enable more compact ACLs by allowing administrators to group hosts with common access control policies into a common IP subnet. For example, campus 3 identifies user machines through a small number of IP prefixes (corresponding to the faculty and student VLANs), allowing concise ACLs for traffic sent by user machines (e.g., to ensure only SMTP traffic is allowed to reach the email servers on the infrastructure VLAN).

Preventing Source IP Address Spoofing — Source IP address spoofing is a serious security problem, since spoofing allows attackers to evade detection or shift blame for their attacks to others. Assigning host addresses from a common IP prefix simplifies the preventive filtering of packets with spoofed source IP addresses. Hosts in the same VLAN are assigned IP addresses from the same subnet(s). This allows network administrators to configure ACLs at the VLAN's gateway router to drop any packets with source IP addresses from other prefixes. Campus 3 does precisely that.

Supporting Quality of Service — Classifying packets based on IP prefixes applies not only to access control, but also to quality of service (QoS) policies. For example, administrators can configure a router to place IP packets in different queues (with different priority levels) based on the source or destination IP prefix, if hosts are grouped into VLANs based on their QoS requirements. None of the campuses in our study apply these kinds of QoS policies.

Decentralizing Network Management

VLANs allow administrators to delegate some management tasks to individual departments. VLANs also simplify network troubleshooting by allowing an administrator to observe connectivity from any part of the campus simply by trunking a port to a VLAN.

Federated Management — Campus network administrators sometimes assign all hosts in one department to a VLAN, so each department can have its own control over its hosts in different locations on campus while sharing the same physical infrastructure. Some campuses allocate

portions of the VLAN ID space to departments and allow those departments to manage their networks independently. For example, campus 1 has a university-wide IT group and many smaller IT groups. The university-wide group allocates a contiguous block of IP addresses to one VLAN and hands it over to a smaller IT group. One IT group manages a "classroom" VLAN that consists of a computer in each classroom across 60 buildings. Campus 2 allocates a portion of the VLAN ID space to the computer science department and provides a web interface to help the administrators manage the router and firewall settings between the department and the rest of the campus. Campus 4 assigns different gymnasiums across the campus to the same VLAN; administrators for that VLAN can then set firewall rules independently from the rest of the campus.

Easier Troubleshooting — VLANs allow network administrators to group hosts based on policy requirements, independent of their locations. If two hosts in the same policy group are in different locations on the campus, administrators can still assign them to the same VLAN so that they can communicate with each other, without interference from intermediate firewalls or routers. In campus 4, the dormitory VLAN spans the campus, including places outside the dormitories; such a setup allows network administrators to help student users diagnose problems since they can put a host on this VLAN anywhere on the campus. Campus 2 also has some VLANs across campus, such as a network-wide VLAN for the IT support team and a VLAN for deploying new experimental management architectures based on OpenFlow [4].

Enabling Host Mobility

VLANs make host mobility easier on a campus wireless network, because hosts can retain their original IP addresses when they move from one access point to another. Allocating a single VLAN to the campus wireless network, as is done in campus 2, allows devices to move anywhere on the campus without having to obtain a new IP address. The campus 2 wireless network has about 6000 active hosts on the same VLAN. These hosts include laptops, mobile phones, passenger counters, and vehicle locators. As users move across the campus on foot or in vehicles, they can remain connected to the campus network, migrating between access points without experiencing disruptions to ongoing connections.

Problem: Limited Granularity of Policy

VLANs are a relatively inflexible way to support policies. In this section, we discuss three main limitations VLANs impose on the granularity of policies — limits on the number of VLANs, limits on the number of hosts per VLAN, and the difficulty of assigning an access port to multiple VLANs without end-host support. We also discuss the incomplete ways administrators try to work around these limitations.

Limited Number of VLANs

The total number of VLANs is limited because of built-in protocol limitations (i.e., VLAN ID space) and implementation limitations (i.e., switch and router resources):

- **VLAN ID space:** The VLAN ID is a 12-bit header field, limiting a network to 4096 VLANs.[1]
- **Switch memory:** Limited memory for storing bridge tables often restricts individual switches to supporting 300–500 VLANs.
- **Router resources:** Inter-VLAN traffic imposes additional load on the routers.

Administrators work around these limitations in two ways.

Placing Multiple Groups in the Same VLAN — Administrators can assign multiple groups of hosts to a single VLAN and configure finer-grained access control policies at the routers to differentiate between hosts in different groups. Campus 1 combines some groups of hosts together, assigning each group a different block of IP addresses within a larger shared subnet. From the configuration data, we see that about 11 percent of the VLANs have ACLs expressed on smaller IP address blocks. For example, one VLAN contains the DNS servers, logging and management servers, and some dorm network web servers. Although these hosts reside in different locations, are used for different purposes, and have different reachability policies, they are placed in a single VLAN because they are managed by an IT group that has a single VLAN ID and one IP subnet.

Reusing the Limited VLAN Identifiers — To deal with limitations on the number of VLAN IDs, administrators can use the same VLAN ID for multiple VLANs, as long as the VLANs do not have any links or switches in common. Unfortunately, reusing VLAN IDs makes configuration more difficult, since administrators must take care that these VLANs remain disjoint as new hosts, links, and switches are added to the network. Campus 1, in particular, reuses VLAN IDs quite extensively.

Limited Number of Hosts per VLAN

The overheads of broadcast traffic, flooding, and spanning tree impose limits on the number of hosts in each VLAN. For example, campus 1 has a wireless VLAN with 3000 access points and thousands of mobile hosts that receive a large amount of broadcast traffic. These scalability limitations make it difficult to represent large groups with a single VLAN. Administrators work around this problem by artificially partitioning these larger groups.

Dividing a Large Group into Multiple VLANs — A large group can be divided into multiple VLANs. For example, campus 1 has public computer laboratories with 2500 hosts across *16* VLANs. The 1200 hosts in one academic college in campus 1 are divided into *eight* VLANs. Dividing a large group into multiple VLANs unfortunately prevents mobile hosts from retaining their IP addresses as they move

VLANs allow administrators to delegate some management tasks to individual departments. VLANs also simplify network troubleshooting by allowing an administrator to observe connectivity from any part of the campus simply by trunking a port to a VLAN.

[1] *IEEE 802.1QinQ provides a way to extend the ID space using multiple tags.*

To deal with limitations on the number of VLAN IDs, administrators can use the same VLAN ID for multiple VLANs, as long as the VLANs do not have any links or switches in common. Unfortunately, reusing VLAN IDs makes configuration more difficult.

from one location to another. Additionally, the VLANs must be configured with the same access control policy to retain the semantics that would exist if hosts belonged to a single larger group.

Coarse-Grained Assignment of Traffic to VLANs

Although they are natural for grouping traffic by *end host*, VLANs are a clumsy way to group traffic across other dimensions (e.g., by application). With end-host support for VLAN tagging, hosts can assign different virtual interfaces to different VLANs. For example, a computer hosting multiple virtual machines can run a software switch that has a different access port (and, hence, can assign a different VLAN) for each virtual interface. However, the end host must support VLANs making it hard to work with the heterogeneous user devices common on college campuses. In addition, the campus administrator must *trust* the user machine to faithfully apply the appropriate VLAN tag — introducing potential security risks. Although protocols like 802.1x can help authenticate hosts, many campuses do not force all hosts to use these mechanisms.

Unexpected problems can arise when administrators assign VLANs directly to access ports. For example, campus 3 assigns each access port to a (single) VLAN dynamically, based on the source MAC address of the attached host. If multiple hosts connect to a single wall jack (e.g., via a common hub or an unmanaged switch), the hosts are assigned to the same VLAN — based on the MAC address of whatever host sends the *first* packet. Since campus 3 has different VLANs for faculty and students, this can raise security problems when a student plugs into a hub in a faculty member's office or vice versa. The same problem arises if a single computer runs multiple virtual machines, each with its own virtual interface and MAC address. By connecting to the same switch access port, all of these virtual interfaces would be assigned to the same VLAN, a problem raised by the administrators in campus 2.

Restricting each access port to a single VLAN significantly limits the kinds of policies the network can support. For example, administrators cannot assign a single host interface to multiple groups (e.g., a faculty member in the systems group cannot belong to both the faculty VLAN and the systemsgroup VLAN) or have different applications belong to different groups (e.g., web traffic cannot belong to a different VLAN than Skype traffic).

Problem: Complex Configuration

Although Ethernet was designed with the goal of "zero configuration," VLAN configuration is challenging and errorprone [5], for two main reasons. First, each host's IP address must be consistent with the IP subnet of its VLAN. Second, the switches require configuration to ensure each VLAN has an efficient spanning tree that remains connected under common failure scenarios.

Host Address Assignment

Administrators associate each VLAN with one or more IP subnets and must ensure that the host interfaces within that VLAN are assigned addresses from that block. The tight coupling between VLANs and IP address assignment leads to two problems.

Wasting IP Addresses — All four campuses have a one-to-one mapping between an IP subnet and a VLAN. Since IP prefixes must align with *power-of-two* boundaries, VLANs can lead to fragmentation of the available address space — especially if 5 some VLANs have fewer hosts than others.[2] Campus 1, for instance, originally assigned a /24 prefix to each VLAN but, after running out of address space, was forced to use smaller subnets for some VLANs.

Complex Host Address Assignment — To ensure that host IP addresses are consistent with the VLAN subnets, Campus 1 manually configures each host with a static IP address from the appropriate VLAN, except for a few VLANs (e.g., the wireless network) that use DHCP. The other campuses use DHCP to automatically assign IP addresses based on the hosts' MAC addresses. However, the administrators must ensure that DHCP requests reach the DHCP server, even though broadcast traffic only reaches machines in the same VLAN. Rather than devote a DHCP server to each VLAN, campuses 2, 3, and 4 use *relay agents* to forward requests to a common DHCP server, requiring additional configuration on the routers [6]. Either way, the DHCP server configuration must be consistent with whatever system is used to assign hosts to VLANs.

Spanning Tree Computation

Switches must be configured to know which VLANs they should support on each trunk link. Administrators must explicitly configure both ends of every trunk link with the list of VLANs to participate in. For example, in Fig. 1, VLAN1 must be allowed on the link between S1 and S2, while VLAN2 need not be permitted. Wrongly omitting a VLAN from that list disrupts communication between the hosts on that VLAN. Unnecessarily including extra VLANs leads to extra broadcast/flooding traffic and larger bridge tables. Determining which links should participate in a VLAN, and which switch should serve as the root bridge of the spanning tree, is often difficult.

Limitations of Automated Trunk Configuration — Manual configuration of trunk links is error-prone [7], and inconsistencies often arise as the network evolves [8]. Automated tools, like Cisco's VLAN Trunk Protocol (VTP) [2], reduce the need for manual trunk configuration. However, these tools require administrators to divide the network into VTP domains, where switches in the same domain cooperate to identify which VLANs each link should support. Each switch must participate in all VLANs in its domain, leading to extra overhead; in fact, some commercial switches can only participate in a handful of

[2] *IPv6 might solve the problem but will not be widely deployed in the foreseeable future.*

spanning-tree instances, limiting the effective size of VTP domains. As a result, campus 1 is divided into several smaller VTP domains, using manually-configured trunk links to interconnect the domains. Campus 2 does not use VTP because some of its switches come from another vendor that does not support Cisco's proprietary protocol. Campus 3 does not use VTP because the administrators prefer to know *by design* which links participate in each VLAN, to simplify network troubleshooting.

Enabling Extra Links to Survive Failures — Although Ethernet switches can compute a spanning tree automatically, administrators must often intervene to ensure that each VLAN remains connected after a failure. To prevent partitioning of the VLANs, campus 1 installs parallel links between switches and treats them as one logical link; this ensures that the VLANs remain connected even if a physical link fails. To survive switch failures, campus 1 configures the trunk links between the core switches to participate in all VLANs. In general, identifying which links to include is challenging, since enabling too many links in the VLAN is wasteful, but having too few can lead to partitions during failures.

Distributing Load Over the Root Bridges — The switches near the root of a spanning tree must carry a large amount of traffic. Dividing the network into multiple VLANs can help distribute the load over multiple spanning trees with different root bridges. By default, the switch with the smallest identifier becomes the root of the spanning tree, resulting in the same switch serving as the root bridge in multiple VLANs. To distribute traffic load more evenly, administrators often configure the root bridge of each VLAN manually. For example, the administrators of campus 1 select the most powerful switches to serve as root bridges.

Conclusion

We have surveyed four campus networks to better understand and illustrate how VLANs are used in practice. Our analysis indicates that VLANs are used for many objectives that they were not originally intended for, and are often ill-suited for the tasks Further, the use of VLANs complicates network configuration management. We believe future enterprise networks should look at ways to minimize the use of VLANs and explore more direct ways to achieve the network administrators' objectives with the goal to make management easier for campus and enterprise administrators.

To extend our understanding of the VLAN usage in practice, we call for operators of campus and enterprise networks to participate in the survey available at [9].

Acknowledgments

We thank Russ Clark (Georgia Tech), Brad Devine (Purdue), Duane Kyburz (Purdue), Peter Olenick (Princeton), and Chris Tengi (Princeton) for sharing their expertise and experiences about network management and VLANs.

References

[1] "VLAN Management Policy Server," http://www.cisco.com/en/US/tech/tk389/tk689/technologiestechnote09186a00800c4548.shtml.
[2] "VLAN Trunking Protocol," http://www.cisco.com/en/US/tech/tk389/tk689/technologiestechnote09186a0080094c52.shtml.
[3] "Multiple VLAN Registration Protocol," http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/catos/8.x/configuration/guide/mvrp.pdf.
[4] N. McKeown *et al.*, "OpenFlow: Enabling Innovation in Campus Networks," *ACM Comp. Commun. Rev.*, Apr. 2008.
[5] T. Benson, A. Akella, and D. Maltz, "Unraveling the Complexity of Network Management," *Proc. NSDI*, Apr. 2009.
[6] C. J. Tengi *et al.*, "autoMAC: A Tool for Automating Network Moves, Adds, and Changes," *Proc. Large Installation Sys. Admin. Conf.*, 2004.
[7] P. Garimella *et al.*, "Characterizing VLAN Usage in an Operational Network," *Proc. Wksp. Internet Network Mgmt.*, Aug. 2007.
[8] X. Sun *et al.*, "A Systematic Approach for Evolving VLAN Design," *IEEE INFOCOM*, 2010.
[9] http://www.surveymonkey.com/s/X5K5GLM.

Biographies

Minlan Yu (minlanyu@cs.princeton.edu) is a Ph.D. student in the computer science department at Princeton University. She received her B.S. in computer science and mathematics from Peking University in 2006 and her M.S. in computer science from Princeton University in 2008. She has interned at Bell Labs, AT&T Labs Research, and Microsoft. Her research interest is in network virtualization, and enterprise and data center networks.

Xin Sun is a Ph.D. candidate in the School of Electrical and Computer Engineering at Purdue University, West Lafayette, Indiana, where he works with Prof. Sanjay Rao. His research interests are in the design and configuration of large-scale enterprise networks, and the migration of such networks to new architectures. He received his B.Eng. degree in computer engineering from the University of Science and Technology of China in 2005.

Nick Feamster is an associate professor in the College of Computing at Georgia Tech. He received his Ph.D. in computer science from MIT in 2005, and his S.B. and M.Eng. degrees in electrical engineering and computer science from MIT in 2000 and 2001, respectively. His research focuses on many aspects of computer networking and networked systems, including the design, measurement, and analysis of network routing protocols, network operations and security, and anonymous communication systems. In December 2008 he received the Presidential Early Career Award for Scientists and Engineers (PECASE) for his contributions to cybersecurity, notably spam filtering. His honors include the Technology Review 35 "Top Young Innovators Under 35" award, a Sloan Research Fellowship, the NSF CAREER award, the IBM Faculty Fellowship, and award papers at SIGCOMM 2006 (network-level behavior of spammers), NSDI 2005 (fault detection in router configuration), Usenix Security 2002 (circumventing web censorship using Infranet), and Usenix Security 2001 (web cookie analysis).

Sanjay Rao is an assistant pProfessor in the School of Electrical and Computer Engineering, Purdue University. He obtained his Ph.D. in computer science from Carnegie Mellon University. His current research interests are in enterprise management and cloud computing. In the past, he has done pioneering work on live streaming using peer-to-peer systems. He is a recipient of an NSF Career award, and has served as a Technical Program Chair of the INM/WREN workshop.

Jennifer Rexford is a professor in the Computer Science Department at Princeton University. From 1996 to 2004 she was a member of the Network Management and Performance Department at AT&T Labs–Research. She is co-author of the book *Web Protocols and Practice* (Addison-Wesley, May 2001). She received her B.S.E. degree in electrical engineering from Princeton University in 1991 and her Ph.D. degree in EECS from the University of Michigan in 1996.

We believe future enterprise networks should look at ways to minimize the use of VLANs and explore more direct ways to achieve the network administrators' objectives with the goal to make management easier for campus and enterprise administrators.