Implementation of Virtual LANs over ATM WANs

Torsten Braun[a] and Martin Mähler[b]

[a] Institute of Computer Science and Applied Mathematics, University of Berne
Neubrückstr. 10, CH-3012 Berne
Phone: +41 31 631 4994, Fax: +41 31 631 3965, Email: braun@iam.unibe.ch

[b] IBM European Networking Center
Vangerowstr. 18, D-69115 Heidelberg
Phone: +49 6221 59 4406, Fax: +49 6221 59 3300, Email: maehler@heidelbg.ibm.com

ABSTRACT

Virtual LANs (VLANs) allow to interconnect users over campus or wide area networks and gives the users the impression as they would be connected to the same local area network (LAN). The implementation of VLANs is based on ATM Forum's LAN Emulation and LAN/ATM switches providing interconnection of emulated LANs over ATM and the LAN ports to which the user's end systems are attached to. The paper discusses possible implementation architectures and describes advanced features such as ATM short-cuts, QoS, and redundancy concepts.

Keywords: virtual LANs, LAN Emulation, Switching, ATM

1. INTRODUCTION

Local Area Networks (LANs) such as Ethernet or Token Ring are widely used to interconnect computers in companies, universities or other institutions. LANs are the basis for different protocols such as TCP/IP, SNA, Netbios, Appletalk etc. LANs have initially been implemented as shared media networks, i.e. a relatively high number of end systems are connected to a common media and have to share the available network bandwidth.

During the last years, LAN switches have been introduced providing an increased number of ports with a small number of connected nodes or with even single nodes. The LAN switches perform bridging/switching functions to forward packets among the ports. The concept of connecting a few nodes to a single port (micro-segmentation) provides several benefits. Less nodes have to share the available bandwidth of the transmission medium. Performance problems due to media access collisions are reduced. Ports can be grouped together according to their logical working group membership or to their layer-3-subnetwork address. Reconfiguration does not require cabling changes but only switch management and due to the hierarchical cabling structure, cable disruptions have only effect for the nodes connected to the corresponding port.

The next logical evolution step of micro-segmentation are virtual LANs (VLANs). A VLAN is a broadcast domain (e.g., an IP subnet) that is independent of the physical topology with membership governed by a set of policies or rules. Switching or bridging is used within VLANs and routing is used between VLANs. All devices of a VLAN such as bridges, routers, and end systems have the same view as if they would be connected to a common regular LAN. VLANs are a technology for companies willing to interconnect different LAN attached endsystems (PCs, workstations, servers) via a backbone network. The backbone network may consist of a local ATM network to support VLANs over a campus or a public ATM network to support VLANs over the wide area.

VLANs provide all the LAN benefits over public wide area ATM networks. They are implemented based on LAN switches which are interconnected over ATM using LAN emulation (LANE) [1]. Campus ATM switches are connected to public wide area ATM switches on one side and to LAN/ATM switches on the other side. In addition, network equipment providing services for LANE (LECS, LES BUS), bridging, routing etc. are connected to the campus ATM switches too.

2. VLAN BENEFITS

LANs are mainly required within company networks covering a single location or campus. However, companies are expanding and branches are more and more distributed over different locations. More often, users at different locations have to work together and have to share data located at common servers. Such a scenario, which occurs with a single but distributed company or with a so-called virtual company, would highly benefit from VLANs. The ATM-based VLAN technology allows to meet emerging customer requirements that can not be met by legacy LANs such as high bandwidth and other properties required for multi-media communication. The VLAN technology provides several benefits for network providers and the customers' network administrators, operators, and users.

- Flexibility
  Future networks must support logical structures and organizations within a company and between companies such as logical workgroups. Workgroups are small or large collections of end-users sharing computing resources. Creation of workgroups has to be supported irrespective of the physical infrastructure, i.e. workgroups have to be defined based on attributes of users and their end systems such as IP subnet addresses, even if workgroup members are geographically distributed or belong to different companies.
- Support of Workstation / User Movements within a Company
  Today, company organizations are becoming more and more dynamic. The mobility and movements of users will cause considerable expense associated with re-wiring and network management tasks. Thus, administration of moves, adds, and changes of devices within a building or campus has to be supported. In a VLAN environment, host moves, adds, and changes can be implemented automatically since VLAN membership is independent of physical topology. If a user or a server moves from one location to the other, no migration concepts or expensive re-wiring must be performed. The old port of the user's end system or the server must be removed from the VLAN, while the new port has to be added to the VLAN. These changes can be made effective by software re-configuration. If a certain set of ports has already been assigned to a VLAN, the users can arbitrarily move between these ports even without any re-configuration. In addition to port-based VLANs, several sub-VLANs can be built within the VLAN based on MAC addresses, protocol IDs, IP subnet addresses, IP multicast group address, user-defined protocol header fields etc. A port is assigned to a sub-VLAN based on traffic analysis on a given port. For example, if an IP and an IPX sub-VLAN have been defined, the access devices learn which kind of protocol is running at the attached end systems. Some ports are assigned to the IP VLAN, while others are assigned to the IPX VLAN, some may be even assigned to both. In that case, a pure IP port does not see any IPX frame. This avoids unnecessary flooding of broadcast traffic.
- Mobile Users
  Portable computers are becoming more and more popular. For example, sales persons might want to access the large databases in their companies while negotiating a contract with their customers. Mobile user support has to enable mobile users to get access to the network from any point without changing configuration or addresses on their computer. Mobile users simple log in from any point over public networks or over the Internet and participate in their VLAN environment as sitting in their own office.
- Performance
  Performance provided by the ATM-based VLAN technology is comparable to that of a regular LAN allowing to run any application behaving well in the local environment also over the wide area. Popular examples are client/server applications, business applications based on SNA, Intranet and Internet applications, multimedia applications etc. The bandwidth available to a user is limited by the bandwidth of the access link, i.e. the link between the end system and the access switch, and the bandwidth of the ATM network interconnecting the switches. However, ATM is a scalable technology with available products supporting access bandwidths up to 622 Mbps today. Moreover, delay characteristics of ATM are much better than for legacy LAN technology. Broadcast traffic can be reduced by intelligent broadcast avoidance mechanisms residing in LAN/ATM switches. This helps to decrease the number of MAC frames sent over a port to the necessary minimum.
- Security
  The VLAN concept is an ideal basis to provide closed user group services. To implement such a closed user group, one has to ensure that all data can be accessed by members of a specific group. An important security issue is limitation of broad-

cast traffic to VLANs. Data sent within a VLAN must not be broadcasted to devices that are not member of the VLAN in order to prevent unauthorized devices to monitor traffic passing the borders of a VLAN. Only the hosts of a common VLAN receive broadcast and multicast frames originating within the common VLAN. Security of VLANs is similar to the security level of physically separated networks. This allows to avoid the implementation of costly security mechanisms such as password administration for server access, firewalls etc.

- Server Centrex
  LAN attached end systems may belong to as many VLANs as they have network interfaces. The ATM technology enables the creation of virtual network interfaces (LAN Emulation clients, LECs) allowing the end system to assign each virtual network interface to a different VLAN. An ATM attached server system can be directly attached to all the VLANs to which the server-specific services shall be provided. ATM attached servers, end systems or LAN switches can be connected via direct ATM connections. In the case of geographically distributed VLANs, single servers can serve all the users of the VLAN independent of their location. It is, therefore, not necessary for a company to run replicated servers at each of their locations. A network provider can connect servers offering value-added services such as WWW, file, e-mail, backup, or other application-specific servers directly to the public ATM network in order to enable high-speed server access from the users' end systems.
- Quality of Service
  Multimedia and real-time applications like video conferencing or telephony require dedicated bandwidth for high-speed data transfer that is not slowed down by high congestion and latency. In contrast to other network technologies, ATM is able to guarantee Quality of Service (QoS) over the local or the wide area backbone network. Dedicated ATM connections can be used to interconnect the LAN segments which form the VLAN. ATM backbones avoid performance degradation as they happen in conventional connectionless networks such as router networks.
- Multiprotocol Support
  VLANs provide LAN services such as Ethernet or Token Ring. A LAN service is independent of higher-layer protocols. All commonly used protocols such as TCP/IP, SNA, DECnet, AppleTalk and even non-routeable protocols such as Netbios can run on top of a LAN or VLAN service. This allows the VLAN technology to be used in arbitrary network environments. VLANs can be the basic platform to implement IP networks such as Internet subnetworks, Intranets or Extranets. Companies with SNA equipment can establish SNA-specific VLANs transporting SNA traffic directly on top of the LAN service without using encapsulation techniques such as data link switching. Customers with mixed traffic can establish several individual VLANs, each VLAN for a certain protocol family. For example, a company with IP and SNA traffic can establish an IP VLAN and an SNA VLAN, both VLANs on top of the same network infrastructure.
- Costs
  The public VLAN service allows significant cost reduction for customers due to several reasons. First, the VLAN service allows a customer to implement secure networks very easily without installing and administering firewalls, password based authentication procedures required for controlled server access, etc. Second, servers need not to be replicated for several locations if a company is distributed over many locations. For example, a company with N locations does not need to install and run N replicated servers but only a single server for each server function (e.g., mail, backup, file services). The maintenance of such servers in a single server farm is much cheaper than supporting each server function at each location. Another reason for reduced costs is outsourcing of LAN management. LAN management tasks can be provided by the VLAN service provider. The VLAN service provider can support all the customers by a single management center. The management costs for the provider are, therefore, cheaper than the sum of all the customer's management costs in the case the customers manage their networks by themselves.

3. VLAN IMPLEMENTATION

3.1 Implementation Architecture

The following Figures show the implementation of the VLAN service based on IBM 8260 campus ATM switches, IBM 8273/ 8274 LAN/ATM switches and Multiprotocol Switched Services (MSS) servers. A MSS server provides several different ser-

vices for LANE (LAN Emulation Server, LAN Emulation Configuration Server), classical IP over ATM (ATM ARP Server) [5], bridging, and routing (IPX/IP routers, interior and exterior routing protocols, NHRP).

The campus ATM switches located at the different sites of the VLAN network are interconnected via permanent virtual paths (PVPs) provided by the ATM WAN. Many ATM WAN providers currently do only support PVPs but no SVCs. The campus ATM switch terminates the PVP and maps all SVCs required for LAN Emulation to the PVPs. The campus ATM switch connects the LAN/ATM switches and other ATM devices (LAN emulation servers, LAN emulation configuration servers, routers etc.) to the ATM network. This enables the ATM devices to establish SVCs transparently over the ATM WAN. The LAN/ATM switches build the interface between LAN-attached end systems or other devices such as routers and the ATM network. At each LAN/ATM switch members of different VLANs may be connected to the LAN (Ethernet, Fast Ethernet, Token Ring, FDDI) ports.

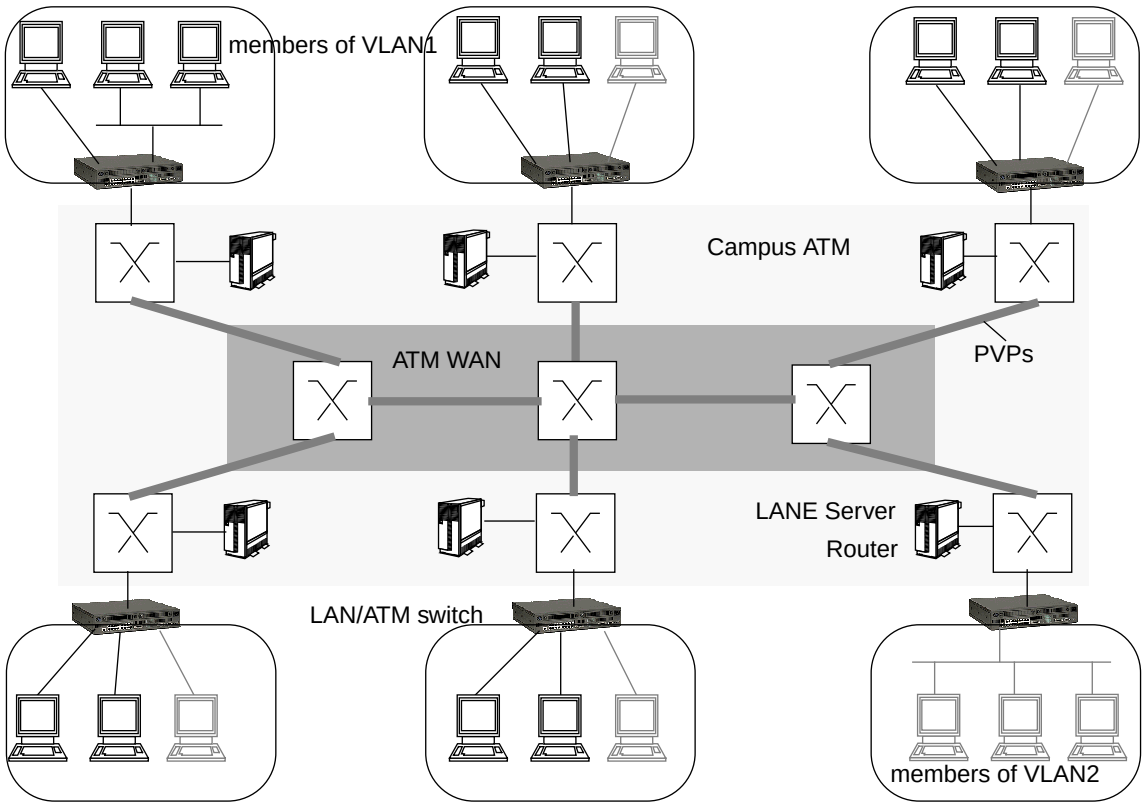

Figure 1: VLAN Implementation over ATM WANs

The VLAN network architecture is shown in Figure 2. Each VLAN consists of an ELAN as the backbone network and a set of LAN ports assigned to the VLAN. In Figure 2, two VLANs, namely VLAN1 (black) and VLAN2 (grey) are realized.The LAN/ATM switches are connected to the ELAN forming the VLAN backbones. E.g., ELAN1 is the backbone for VLAN1 and ELAN2 is the backbone for VLAN2.

The example scenario consists of four sites interconnected to each other via an ATM WAN. One or more LAN/ATM switches may be installed at each site in order to provide VLAN access for LAN-attached clients (PCs/workstations). Several LAN/ATM switches of a single site are interconnected by local ELANs (ELAN1a, ELAN1b, ELAN2a, ELAN2b) in addition to the global ELANs (ELAN1, ELAN2). This allows that in error cases, when the global ELANs are not accessible, e.g. due to a LAN Emulation Server (LES) failure, the LAN/ATM switches at a single site can still reach each other via the local ELANs. The bridge parameters have to be set such that the spanning tree algorithm enables the bridge ports to the gloabl ELANs.

Servers can be connected directly to the ATM network. This allows to implement a central server farm (server centrex) offering

services to be used by the individual VLAN clients. In Figure 2, the left server offers services to VLAN1, while the other one offers services for VLAN2. Examples for those services are file, WWW, backup, document, video, mail servers or proxy servers providing access to the Internet.

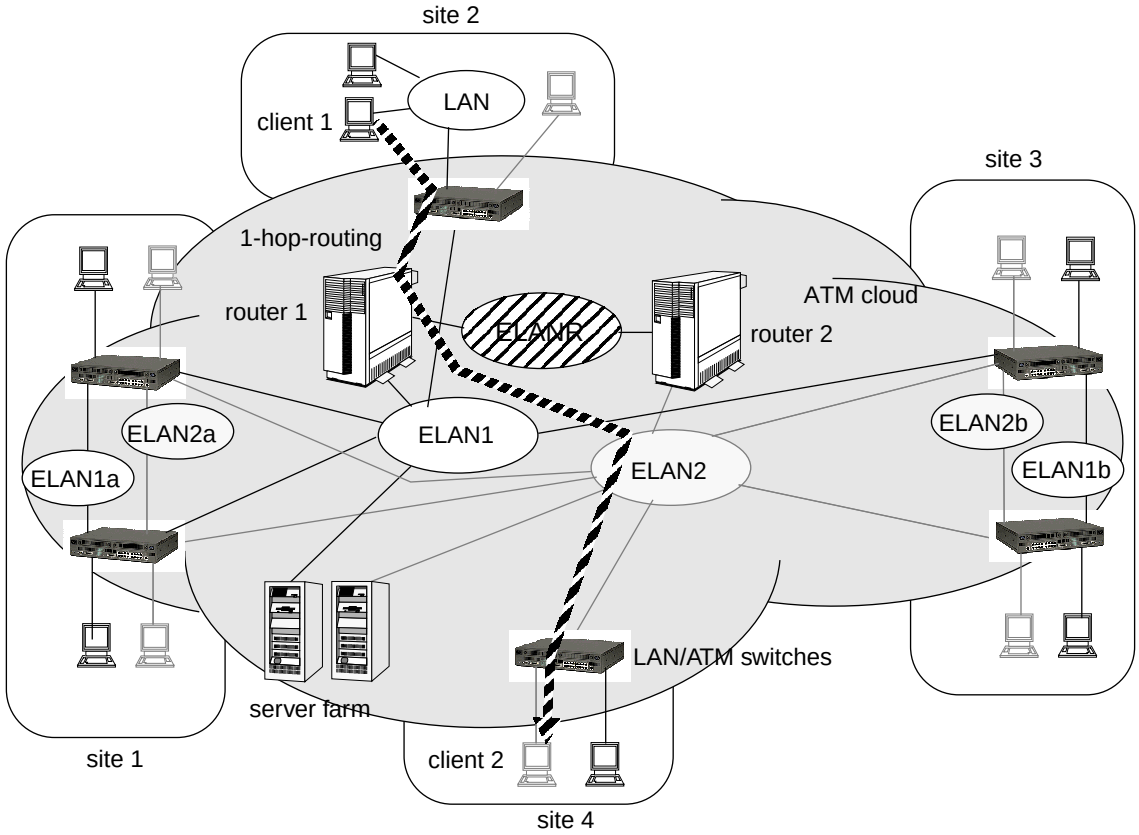

Figure 2: VLAN architecture

3.2 Interconnection of VLANs

VLANs may be used to implement closed user groups. In that case, only VLAN members can reach each other. If required, interconnection of VLANs may be performed via routers. A limited set of VLANs can be interconnected using routers and installing address filters on the routers in order to allow communication among certain VLANs and to prevent communication among others. In Figure 2, each VLAN has its default router. All routers are connected via special router ELANs that are used only for router interconnection, e.g. ELANR. The routers may then run routing protocols such as RIP, OSPF, BGP, etc. over these dedicated router ELANs.

If a client of VLAN1 (client 1) wants to communicate with a client of VLAN2 (client 2), the packet is sent from client 1 via ELAN1 to router 1, via ELANR to router 2, and finally via ELAN2 to client 2. For the implementation of routers, MSS server release 1.1 [3], has been used. MSSr1.1 supports all major routing protocols and supports NHRP [4] over classical IP and LANE networks. In our scenario, using MSSr1.1 establishes a short-cut, i.e. a direct ATM connection between router 1 and the LAN/ATM switch to which client 2 is connected. MSSr1.1 not only supports short-cuts in classical IP networks but provides proprietary NHRP extensions in order to support short-cuts over LANE networks.

3.3 Quality-of-Service

QoS is supported by the MSS server similarly as the QoS support defined in LANEv2 [2]. LECs can register the QoS parameters they can support at the LES. Requesting LECs receive the QoS parameters with an LE_ARP_RESPONSE from the LES.

The LECs of an MSS server can use these QoS parameters in order to setup reserved bandwidth ATM VCs.

The QoS parameters for all ELANs can be configured centrally at the LECS. A LEC retrieves the LES address and the QoS parameters to be registered at the corresponding LES from the LECS. This allows to configure certain QoS values for different ELANs. QoS parameters can also be configured for each LEC individually. However, this requires to set-up each LEC appropriately.

3.4 Redundancy

MSSr1.1 provides all the LANE services such as LANE server (LES), broadcast and unknown server (BUS), LANE configuration server (LECS) and a classical IP ARP server. A critical issue in LANE and classical IP networks is the reliability of the different services. MSSr1.1 allows to establish redundant LESs and ARP servers. In both cases, the backup server (either the LES or the ARP server) establishes a control ATM connection to the primary server in order to check whether the primary is operational or not. If the backup server detects a failure of the primary one, it takes over the primary server function.

In the case of LES/BUS failures, a LEC asks the LECS for the backup LES/BUS. In the case of ARP server failures, the backup ARP server registers the primary ARP server's ATM address at the ATM switch and receives all the ARP requests.

LECS redundancy is achieved by ILMI. In the case of a LECS failure, a LEC takes the next LECS in the list retrieved by ILMI from the ATM switch. Therefore, both primary and backup LECSs, must be configured in the LECS table at an ATM switch.

MSSr1.1 also provides redundancy mechanisms for IP gateways. In that case, both routers have an individual IP/MAC address pair and share a common IP/MAC address pair. Normally, the primary router listens to the shared address pair. When the backup detects a failure of the primary gateway, it begins to listen to the shared address pair and forwards packets sent to the shared MAC address. All clients must in that case use the shared IP address as the default IP gateway address. Redundancy can be combined with load sharing. For example, one MSS server could be used for the primary server functions for VLAN1 and for the backup server functions for VLAN2, while the second MSS server could be configured with primary server functions for VLAN1 and with backup server functions for VLAN2. In that case, the MSS servers share the total load in the normal case, while one MSS server takes over the functions of the other MSS server if that one fails.

3.5 Short-Cut Bridging

A drawback of the architecture depicted in Figure 2 is that all the ELAN control traffic between LECs and LESs may be exchanged via the ATM WAN, in particular when a LES of an ELAN is located at another site than the connected LEC is. For example, two members of VLAN1 located at site 1 communicate via ELAN1 as long as ELAN1 is accessible. ELAN1a is only used for local redundancy when ELAN1 fails. In the normal case, the LAN/ATM bridge ports to ELAN1a should be blocked.

The architecture depicted in Figure 3 avoids the drawbacks of the previous architecture. In that case, the LAN/ATM switches of sites 1 and 3 are connected to local ELANs only. The local ELANs are connected to the global ELANs via so-called short-cut bridges implemented using MSSr2. These short-cut bridges are able to establish short-cuts between LECs belonging to different ELANs. For example, short-cuts between a LAN/ATM switch at site 1 and a LAN/ATM switch of site 3 can be established as long as the ELANs to be interconnected belong to the same VLAN. The short-cut bridging capability is based on forwarding LE_ARP_REQUEST messages between the different ELANs. The LES for ELAN1a must in that case forward such a request towards the LES of ELAN1b that gives the response containing the desired ATM address of the remote LAN/ATM switch. This architecture can also reduce the amount of broadcasts transferred over the ATM WAN, e.g. for IP ARPs. An intelligent broadcast manager running on the short-cut bridge maps specific broadcasts to unicast MAC addresses.

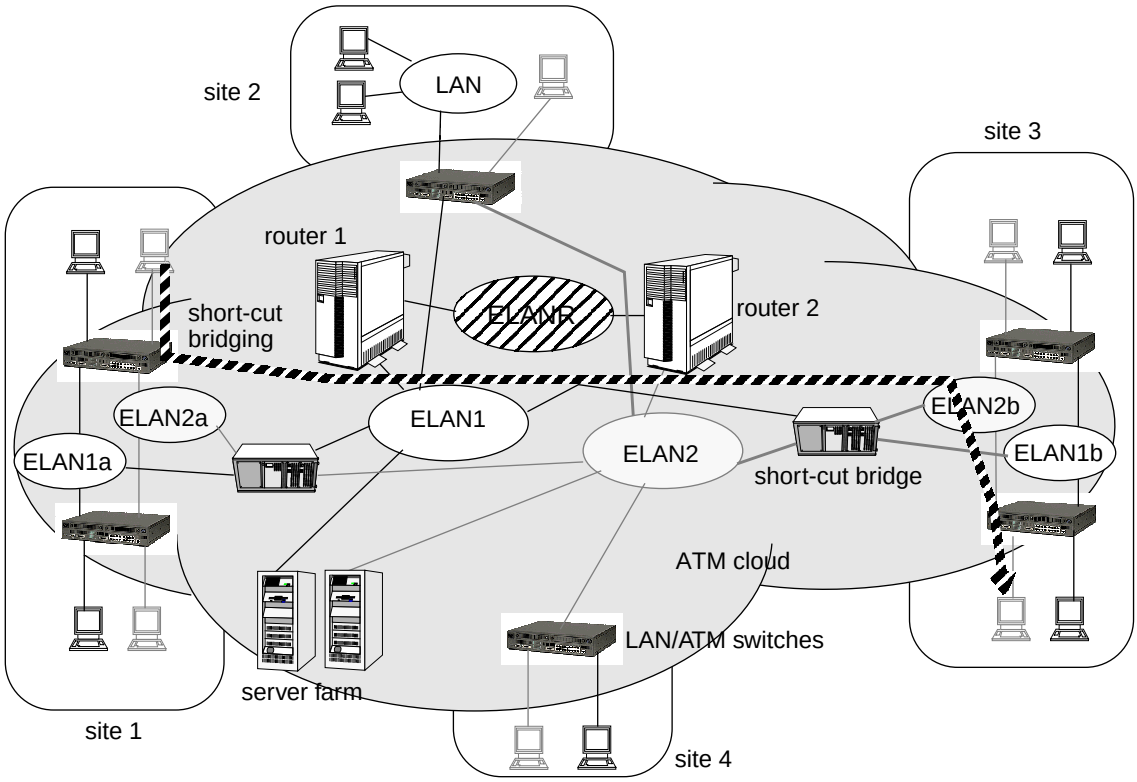

Figure 3: VLAN architecture with short-cut bridging

4. PERFORMANCE

Detailed performance measurements have not yet been performed. The applications that have been used in a wide area scenario performed like in a local LAN environment. Measurements of delays and round-trip times in a scenario where the different sites have been several hundreds of kilometers away from each other showed that the delay depends mainly on the signal propagation delay and that delays due to switching delays by LAN or ATM switches are nearly negligible. This is a significant difference to router networks.

5. CONCLUSION AND OUTLOOK

The VLAN technology presented in this paper is a promising approach for the implementation of various networks, in particular for multi-protocol networks with high performance and security requirements. In addition, costs can be reduced by central server farms and centralizing or even outsourcing network administration and management. In summary, VLANs provide an ideal basis for the implementation of virtual private networks (VPNs), Intranets, and Extranets. IP tunneling protocols such as L2TP can be used to connect users via the Internet to their home VLANs.

However, there remain a few issues to be solved or improved in the future. If a service provider uses the same network equipment for different VLANs of different customers (e.g., if they share the same office building), it would be unwise to allow one customer full network management capabilities over the network equipment. One the other hand, both customers want to have the ability to view or reconfigure their VLAN. A dedicated network management system for these requirements is necessary.

Another issue is security, in particular in SVC-based WANs. A LAN switch connected to the ATM WAN usually accepts all incoming connection setup requests. This would allow a third party to establish ATM connections to a LAN switch and to connect to a VLAN. One approach to prevent this are access lists to be configured in the edge ATM switches, e.g.in the campus ATM switches. In multi-provider networks stronger mechanisms based on authentication mechanisms may be required.

6. REFERENCES

[1] ATM Forum Technical Committee: LAN Emulation over ATM, Version 1.0, January 1995

[2] ATM Forum Technical Committee: LAN Emulation over ATM Version 2, LUNI Specification, February 1997

[3] C. A. Alexander: A Quick Guide to the IBM Multiprotocol Switched Services (MSS) Server Release 1.1, ftp://ftp.networking.ibm.com/nswww/tr2/tr292260.ps

[4] D. Katz, D. Piscitello, J. Luciani, B. Cole: NBMA Next Hop Resolution Protocol (NHRP), Internet Draft, work in progress, September 1997

[5] M. Laubach: Classical IP and ARP over ATM, RFC 1577, January 1994