A Firewalling Scheme for Securing MPOA-based Corporate Intranets

Jun Xu and Mukesh Singhal
Department of Computer and Information Science
The Ohio State University
Columbus, OH 43210
{jun,singhal}@cis.ohio-state.edu

Abstract

Enterprise intranets are vulnerable not only to attacks from the Internet barbarians, but also to internal security breaches. Firewalling, typically used to control accesses to a trusted network from the Internet, is also an effective approach to secure enterprise intranets. Firewalling in traditional enterprise intranets is typically implemented as packet filtering routers sitting between different subnets to filter traffic between them. Such firewalling will not work in future ATM-based enterprise intranets that employs MPOA as the internetworking infrastructure. The reason is that in MPOA, cut-through connections intended for high-bandwidth and low-delay end-to-end communication will bypass such packet-filtering routers. So far there is no satisfactory solution to this problem. In this paper, we propose a firewalling scheme for MPOA, which is tight in security, low in cost, high in performance, and easy to manage. Based on the novel concept of "logical chokepoints", the proposed firewalling scheme does not require the existence of physical chokepoints inside the network to be secured. The proposed firewalling scheme is centrally managed so that it scales well to very large MPOA-based enterprise networks.

Key Words: Security, Intranet Firewalling, MPOA, Packet Filtering, Virtual LAN.

1 Introduction

Enterprise intranets are vulnerable not only to attacks from the Internet barbarians, but also to internal security breaches [4]. In an enterprise intranet, we would like to allow maximum sharing of information and resources within a departmental subnet while enforcing tight control over the cross-departmental network accesses. To guard against attacks from the Internet, a packet-filtering firewall is typically deployed at the chokepoint between an enterprise network and the Internet to inspect the packets that travel between them [3]. Similarly, an enterprise intranet can be secured by placing packet-filtering firewalls at the chokepoints between departmental subnets. Such a firewalling scheme works well over traditional enterprise intranets where shared-media LANs, each of which is an IP subnet, are interconnected by IP routers. However, it may not work effectively with some other internetworking paradigms in which boundaries between security domains are blurred, such as *Multi-Protocol Over ATM* (MPOA).

It is envisioned that future enterprise networks will employ more and more ATM technology as traditional shared-media LANs such as Ethernet and Token Ring hit performance and management limitations. However, the acceptance of ATM technology depends largely on its ability to interoperate with legacy LAN and WAN technology due to the huge past investment in the latter [2]. MPOA is proposed as a unified framework that allows MAC (Media Access Control) and internetwork layer protocols (e.g., IP, IPX, AppleTalk) be transparently transported over an ATM network [6]. In MPOA, ATM-attached *LAN switches*, called *edge devices*, are used to connect the legacy LAN segments to the ATM network. *Edge devices* implement LANE (LAN Emulation) protocol [5], which, together with LANE-enabled ATM-attached hosts, provide a *virtual LAN* (VLAN) service, where ports on an *edge device* can be assigned to particular VLANs, independent of the physical location. *Edge devices* that contains one or more ports that belong to a VLAN serve as multi-protocol bridges to connect physically distributed *LAN segments* and ATM-attached hosts

into a single *broadcast domain*, over which an internetwork layer subnet (e.g., IP subnet) is defined. Communication within a *virtual LAN* now occurs across the ATM network using LANE, which takes advantage of ATM's high bandwidth. However, with LANE, inter-VLAN communication still needs to travel through a router, which renders the router a performance bottleneck and incurs a huge one-hop delay. MPOA is proposed to allow *cut-through connections* be established between *edge devices* on behalf of communicating end systems that belong to different VLANs. With regard to *cut-through connections*, MPOA performs internetwork layer routing and forwarding function over ATM using a *distributed router* approach. A router consists of two functions, namely *route calculation* and *packet forwarding*. *Route calculation* is to run the routing protocol (e.g., OSPF, RIP) to exchange next-hop and reachability information with other routers, and update the routing table. *Packet forwarding* is to forward a packet to the next-hop that is specified in the entry corresponding to the destination address in the routing table. In MPOA, only MPOA servers implement *route calculation* whereas *edge devices* only implement *packet forwarding*. To establish a *cut-through connection*, the *ingress edge device*, to which the source end system is attached to, needs to obtain the ATM address of the *egress edge device*, to which the destination end system is attached to, from an MPOA server through MPOA protocol. Once the *cut-through connection* is established, the ingress/egress edge device handles the *packet forwarding* between the LAN port that the source/destination end system is attached to, and the ATM interface. In this way, the functions of a router are distributed to MPOA servers and edge devices, which, by interacting through MPOA protocol, function as a *distributed router* [1]. The main justification to the approach is the cost. *Route calculation* requires fast CPU and large memory to execute the complex routing protocol and store the distance matrix and is therefore expensive to implement; whereas *packet forwarding* employs simple logic (routing table lookup), and therefore can be implemented in inexpensive ASIC (Application Specific Integrated Circuits) chips. Putting inexpensive functions on large number of

edge devices and expensive functions on a small number of MPOA servers makes MPOA extremely cost-effective [1]. Another justification is that centralized route processing facilitates centralized management. With the *distributed router* approach, only a small number of MPOA servers need to be configured and managed [11].

Cut-through connections allow direct inter-VLAN communication over an ATM network without requiring an intervening router. The flip side of the coin is that the traffic inside cut-through connections bypasses packet filtering firewalls if there are any. Here we assume that each VLAN (or each group of VLANs) corresponds to a distinct security domain, which is in line with its purpose to create virtual workgroups [13]. Intra-VLAN communication is considered secure while inter-VLAN traffic needs inspection. Cut-through connections will lead to security breaches if they are not properly controlled. Two firewalling schemes for MPOA have been proposed, but none of them is satisfactory. One approach considered but later refuted by ATM forum is to let each edge device implement packet filtering function [11]. The approach was later refuted owing to its high cost. As discussed above, for MPOA to be cost-effective, edge devices have to be made inexpensive. As full-fledged packet filtering is a costly feature to implement, by making each edge device a "mini-firewall," the cost of the *distributed router* approach becomes prohibitively high. Moreover, for edge devices to perform packet filtering, complex configuration commands need either be programmed into each edge device or downloaded from MPOA servers using yet another new protocol [11], which incurs high cost in system management and protocol development. The second approach, also proposed by ATM forum, argues that to equip MPOA servers with packet filtering capability, MPOA naturally provides a way to secure cut-through connections as follows. In MPOA, the first several packets from an edge device to a destination internetwork layer address are certain to travel through the default path[1]. An edge device allows the establishment of a cut-through

[1]The only exception happens when the edge device receives a MPOA *trigger* from the MPOA server.

connection only when the first several packets of a *flow* withstand the scrutiny [11]. A problem with this approach is that it can not guard against the *spoofing attack*. For example, the source end system can first send some "benign" packets to the destination end system to simply fool the MPOA servers along the *default path*. Once the cut-through connection is established, the source end system can do whatever it likes. Another problem is that a cut-through connection triggered by an end system A to communicate with an end system C may be accessed by another end system B that is attached to the same edge device, if B also communicates with C. In this case, whereas first several packets between A and C has been inspected by the MPOA servers, no packet between B and C is checked. This is obviously a security hole with the second approach.

In this paper, we propose a firewalling scheme for MPOA-based enterprise intranets which is tight in security, low in cost, high in performance, and easy to manage. Whereas traditional firewalling schemes require the existence of physical chokepoints between security domains to place packet-filtering firewalls, the proposed firewalling approach is based on the novel concept of "logical chokepoints" as follows. In the proposed scheme, an extended MPOA protocol is executed at edge devices and MPOA servers to decide whether a cut-through connection is considered safe and force "unsafe" connections go through some points (the logical chokepoints) inside the ATM network, where high-performance ATM packet-filtering devices are deployed. The concept of logical chokepoints are especially suitable for MPOA environment for two reasons. First, in MPOA, the boundary between security domains (virtual LANs) are blurred due to the existence of cut-through connections over the ATM network. Unlike in traditonal TCP/IP networks, a router no longer physically separates two LANs. Second, it is also not desirable to have a physical chokepoint inside the ATM network because it will be the bottleneck on the backbone.

In Section 2, we present the functions and availability of ATM firewall devices. Section 3 details our firewalling scheme and points out its advantages. Section 4 concludes the paper.

2 ATM Firewall Devices

The availability of high-performance ATM firewall devices, which can perform packet-level filtering over an ATM network at the rate no less than OC-3c (155 Mbps), is a prerequisite of our firewalling scheme. A traditional router-based firewall is not suitable for packet-level filtering over an ATM network for two reasons. First, a packet-filtering router needs to terminate end-to-end ATM connections in the middle in order to extract IP packets for inspection. This involves the process of reassembling incoming cells into packets and disassembling outgoing packets into cells, which incurs a huge overhead. Second, the filtering bandwidth of a traditional firewall is generally below 100Mbps [12], which is far less than the typical ATM rate of OC-3c (155 Mbps) and OC-12c (622 Mbps). Therefore, to filter packets over an ATM network, an ATM firewall device which operates at a high throughput and introduces low delay is called for.

At the time of writing, one commercial ATM firewall product, StorageTek's ATLAS, is available [10]. ATLAS is a *line filter* that scans an ATM physical link to perform the packet-level filtering at OC-3c (155 Mbps) and will reach OC-12c (622 Mbps) in the near future. Two performance boosting strategies are used in ATLAS, which correspond to the two factors that render traditional firewalls unsuitable for an ATM environment. First, to avoid reassembling and disassembling of packets, for each packet ATLAS only inspects its first cell, which contains the IP header and TCP ports information to determine whether the packet is "safe." If the packet is considered safe, all following cells that belong to the packet are passed or otherwise dropped. Second, ATLAS uses a *policy cache architecture* [10] to dramatically speed up the process of deciding whether a packet header is safe or not. The core unit of this architecture is a cache block called *policy cache*. Each entry of the *policy cache* is a combination of TCP/IP source/destination, addresses/ports, and TCP option bits (e.g., ACK, SYN) that are considered "safe". When the first cell of a packet arrives, it is compared with each entry in the *policy cache*. If cache hit occurs, all cells that belong to the

packet are forwarded. Otherwise, the first cell will go through a software screening process and other cells will be buffered in a queue. If the cell is found unsafe, all cells that belong to the packet are dropped. Otherwise, all cells that belong to the packet are allowed to pass, and an entry that contains the header pattern of the packet is now added to the *policy cache*. The *policy cache* is implemented using Content Addressable Memory (CAM), which enables simultaneous search of all memory locations to find a match with the pattern that is searched for. Therefore, ATLAS can quickly decide whether a packet header hits the *policy cache*.

As a *line filter*, ALTAS is designed to be placed on the only link (the chokepoint) between an trusted ATM LAN and an untrusted public ATM WAN. However, ATLAS is not suitable for our MPOA firewalling scheme because it does not understand the PNNI signaling protocol. Whenever a new PVC or SVC is established, TCP/IP rules for that VC have to be manually configured. This is unacceptable in MPOA, where a large number of cut-through connections need to be established on-the-fly within a short period of time. A firewall device based on an ATM switch is more desirable as it is capable of generating TCP/IP rules for each new connection. In an ATM firewall device, TCP/IP filtering rules for different connections are different for two reasons. First, from the VPI/VCI of a connection, the firewall device can determine the IP address/prefix of the source ATM endpoint because there is a binding between the IP address/prefix of the source ATM endpoint and its ATM address. The source ATM endpoint is bound to an IP address if it is a host and is bound to an IP prefix if it is a router. Second, a set of security rules can be specified for each legitimate source ATM endpoint, which is translated into what type of TCP/IP services the source ATM endpoint is allowed to access. These two pieces of information are glued into concrete filtering rules to be applied on the new connection [15]. A switch-based firewall device understands the PNNI signaling protocol so that the configuration of the TCP/IP filtering rules for a new connection is naturally embedded into the signaling message processing. Such switched-based

ATM firewall devices are not only theoretically appealing, but also technologically feasible. In [15], we propose an ATM firewall architecture which nicely embeds high-performance packet-filtering mechanisms into a normal ATM switch. It has the following outstanding features:

- Filter TCP/IP traffic over an ATM connection at the rate OC-3c, OC-12c, and higher,
- Automatically generate TCP/IP filtering rules for each new connection , and
- Easy to manage and administer.

Though in this section we mainly addresses the design and availability issues of ATM firewall devices that filter TCP/IP traffic, the same techniques can be applied to the design of firewall devices that filter other internetwork layer protocols such as IPX and AppleTalk. Therefore, in the following discussion we assume the availability of high-performance ATM firewall devices that are capable of filtering multiple internetwork layer protocols, and they are referred to as *firewall switches*.

3 Proposed MPOA Firewalling Scheme

As discussed in the introduction, inter-VLAN traffic follows either the default path or cut-through connections. We assume that MPOA servers perform the packet filtering function so that traffic over default path is properly inspected. The objective of our firewalling scheme is to secure cut-through connections, which consists of three tasks. In the first task, ATM firewall switches are deployed at the "strategic" points of the ATM network of a MPOA system. Two considerations should be accommodated in choosing the location for placing firewall switches inside the ATM network. First, the requirement for a connection to travel through a firewall switch should not significantly increase the average length of a connection path (measured as number of intermediate switches). Second, the firewall switches should be placed in such a way that load balancing among

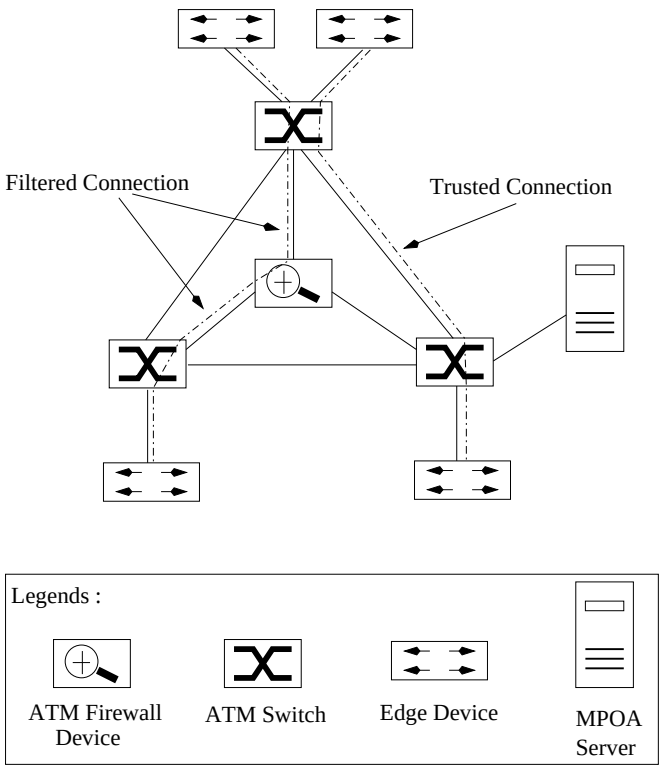

Figure 1: Our firewalling scheme in a sample MPOA system

them can be roughly achieved. The best placement of the firewall switches inside the ATM network depends on the topology and traffic of the MPOA system and can only be determined case by case. The second task, known as *call screening*, is to decide, before a cut-through connection is established, whether the connection is safe. If the connection is found unsafe by *call screening*, the third task is to establish the connection so that it travels through a firewall switch and gets filtered. Such a connection is called a *filtered connection*, as opposite to a *trusted connection*, the content of which needs no inspection. Fig. 1 illustrates our firewalling scheme in a sample MPOA system (oversimplified).

The basic idea of the firewalling scheme is that edge devices act together to form a protection shield against internal security breaches. We assume that in an enterprise network the physical security and configuration of these edge devices will not be tampered with. However, end systems

directly attached to the ATM network will bypass this protection shield. Therefore, our firewalling scheme requires that only trusted servers (e.g., file servers, application servers, mail servers, etc.) are allowed to be directly attached to the ATM switch. All other end systems can only access the ATM network through an edge device. This is not a severe limitation because the-state-of-the-art edge devices can deliver a Quality of Service (QoS) good enough to support high-bandwidth low-latency multimedia applications at a low price.

3.1 Call Screening

In the proposed firewalling scheme, though requests for *call screening* are generated at edge devices where cut-through connections are triggered by *flow detection* process, the call screening decision is not made at edge devices for two reasons. First, a new and complex protocol would be needed for edge devices to download call screening rules from a server and synchronize the local copy with the server's copy from time to time, which is not a trivial task. Second, if *strong authentication* is required for the source end system, the edge device needs also to obtain identity and key information of the end system and perform the authentication [14]. This task is too sophisticated for an edge device to handle.

In the proposed firewalling scheme, the call screening decision is made by MPOA servers. Edge devices interact with MPOA servers to send call screening requests to and obtain call screening replies (decisions) from MPOA servers using an extended MPOA protocol. Here, for ease of discussion, we assume that both the ingress edge device and the egress edge device are servered by the same MPOA server. In case they are servered by different MPOA servers, NHRP flow [9] between MPOA servers can be extended to synchronize information known to both MPOA servers. The MPOA protocol is extended as follows:

1. When edge device A wants to establish a cut-through connection with edge device B, A

will send to the MPOA server a MPOA *resolution request* that contains B's internetwork layer address in order to obtain B's ATM address. For call screening, the format of MPOA *resolution request* will be extended to contain the internetwork layer address of the source end system[2] as well. If *strong authentication* is needed to establish a cut-through connection, authentication information such as a signed digital signature [14] is also contained.

2. When MPOA server receives the MPOA *resolution request*, it makes the call screening decision according to the internetwork layer source and destination addresses contained in the MPOA *resolution request*. The MPOA server also authenticates the source end system according to the authentication information inside the MPOA *resolution request* if applicable. The decision as to whether the connection is secure is contained in the MPOA *resolution reply* from the MPOA server to the edge device.

3. If the connection is determined by the MPOA server as unsafe, the MPOA *cache imposition request* that the MPOA server sends to the egress edge device will solicit for information regarding the distances between the egress switch and each firewall switch (in number of intermediate switches). This information will be used for establishing a *filtered connection* as explained in Section 3.2.

4. When the egress switch receives the MPOA *cache imposition request* from the MPOA server that contains the solicitation, the MPOA *cache imposition reply* will contain the solicited information. This information will also be contained in the MPOA *resolution reply* from the MPOA server to the ingress edge device.

5. When the ingress edge device receives the MPOA *resolution reply* from the MPOA server, it extracts the call screening decision from the reply. Then, if the connection is found unsafe, the

[2]It is actually an optional field of MPOA *resolution request*.

ingress edge device will also extract from the reply the distance information for establishing a filtered connection. Finally, it caches the MPOA *resolution reply* as well as the decision so that future communication between these two end systems may not result in a new MPOA *resolution request*.

In summary, the following new fields are added to MPOA protocol packets:

1. To MPOA *resolution request*: internetwork layer source address, authentication information (optional),
2. To MPOA *cache imposition request*: solicitation for distance information (optional),
3. To MPOA *cache imposition reply*: distance information (optional), and
4. To MPOA *resolution reply*: call screening decision, distance information (optional).

The call screening functions performed at MPOA servers and edge devices result in little increase to their costs. Two call screening functions are performed at MPOA servers: to make call screening decisions and authenticate the source end system. To make screening decisions, MPOA servers need to compare the internetwork layer source and destination addresses contained in the MPOA *resolution request* against a set of *call screening rules*, each of which is a three-tuple <internetwork layer source address, internetwork layer destination address, action>. The *action* field of the entry, in which the first two fields match the internetwork layer addresses of source and destination end systems, specifies whether the cut-through connection needs to be filtered or considered safe. This is obviously an inexpensive table lookup procedure. To authenticate the source end system, the MPOA server needs to search for key information and calculate authentication code such as signed digital signatures [14]. Relatively speaking, this is not an expensive function either. A high-performance cryptographic hardware is available at the price of 2000 dollars [7] each whereas a

full-fledged MPOA server costs at least 20000 dollars. Additionally, this tiny cost increase (about 10%) is further amortized by the large number of edge devices a MPOA server servers.

Call screening functions of an edge device is completely embedded into its MPOA function as follows. To perform *flow detection* and maintain cut-through connections, a normal edge device maintains a table keyed by internetwork layer destination address[3]. The table also contains a *count* field and a *connection* field. If there is a cut-through connection between the ingress edge device and the edge device from which the internetwork layer destination address is reachable, the VPI/VCI of the connection is stored in the *connection* field. Otherwise, the *flow detection* is in progress such that the *count* field stores the number of packets destined for that addresses within a certain time period. In our MPOA firewalling scheme, the internetwork layer destination address alone can no longer be the key. Consider the situation that there are two end systems, B and C, attached to the same edge device. B's traffic to A triggers the establishment of a trusted connection to the edge device to which A is attached. Assume that according to the security policy C is not trusted by A, then C should not be allowed to access the cut-through connection though MPOA may allow it. Therefore, in our firewalling scheme, *flow detection* and MPOA *resolution request* should be initiated for each internetwork layer source and destination address pair, instead of internetwork layer destination address alone. Therefore, the table needs to be extended to contain a new field, the internetwork layer source address. As a consequence, more memory or cache is needed to store the table. The ratio of increase in table size is equal to the average number of end systems attached to the edge device that communicates with the same destination end system. Except for a small number of "hot" servers, the probability for two end system to communicate with the same end system is low. Therefore, we can expect that in a large MPOA system, the ratio of increase will

[3]The key actually includes the identity of a serving MPOA server as well. However, its omission will not affect the validity and performance of our firewalling scheme.

be no larger than two. Since only a small portion of an edge device's memory is used to store the table, and the cost of memory only constitutes a small portion of an edge device's total cost, extra memory required by the table size increase will result in negligible price increase for edge devices.

If the cut-through connection is found safe, the ingress edge device establishes a trusted connection using normal PNNI signaling procedure. Otherwise, the ingress edge device is responsible for establishing a filtered connection, the detail of which is discussed next.

3.2 Establish a Filtered Connection

To establish a filtered connection, the ingress edge device first needs to decide through which firewall switch the connection should travel. It is desirable to choose a firewall switch that will travel through as few intermediate switches as possible. In our MPOA firewalling scheme, each edge device maintains a list of available firewall switches in the ATM network, their ATM addresses, the internetwork layer protocol(s) they filter, and the estimated number of hops between itself and each firewall switch. To calculate the estimated total number of intermediate switches in the filtered connection through a certain firewall switch, the ingress switch needs to obtain the estimated number of intermediate switches between each firewall switch and the egress switch. We have shown in the extended MPOA protocol how such information is obtained from the egress switch. The ingress edge device then perform additions to obtain the total number of hops between itself and each firewall switch.

Once the ingress edge device has so chosen a firewall switch, it establishes a filtered connection as follows. The ingress edge device specifies in the signaling message the ATM address of the firewall switch as the final destination. The true destination is encoded into a *user-to-user information element* or *generic identifier transport information element* [8]. When this signaling message is forwarded to the firewall switch, the firewall switch knows the trick and will retrieve the final

destination from the corresponding *information element* and use it as the new destination to forward the signaling message. In the meantime, the firewall switch set up the packet filtering rules to be applied on the connection as discussed in Section 2. In this way, no extension to PNNI is needed because the above manipulation happens at user level.

3.3 Advantages of the Proposed Firewalling Scheme

The proposed firewalling scheme has following advantages:

- The proposed firewalling scheme is secure and flexible.

  The proposed firewalling scheme is secure because edge devices form a protection shield that forces each and every unsafe connection between different VLANs travel through a firewall switch and get filtered. Unlike traditional intranet firewalls, the proposed scheme does not require the existence of physical chokepoints and therefore does not place any limitations on the topology of MPOA system. This makes the firewalling scheme applicable in any MPOA-based enterprise intranets.

- The cost of the proposed firewalling scheme is low.

  The cost of the proposed firewalling scheme includes the cost to extend the MPOA protocol, the cost to enhance the function of MPOA servers and edge devices, and the cost of firewall switches. The proposed firewall scheme only requires minor modifications to MPOA protocol. No new protocol step is needed and no more than a couple of new fields are added to each protocol packet. Therefore, the cost of extending MPOA protocol is negligible. We have explained in Section 3.1 that the cost associated with functional enhancements to MPOA servers and edge devices is low. The only significant cost involved is the firewall switches. However, since the design of firewall switches is not very complicated [15], the cost to im-

plement it will keep decreasing with the advance of VLSI technology. Additionally, since the number of such firewall switches is small, the cost will be amortized by the huge number of end systems they are able to protect. In conclusion, the total cost to implement the proposed firewalling scheme in a large MPOA system is small.

- The proposed firewalling scheme is easy to manage.

  The management of the firewalling scheme is centralized in the sense that we only need to configure packet-filtering rules and call screening rules on a small number of MPOA servers and firewall switches. Therefore, the security policy is centrally administered except for the synchronization among MPOA servers and ATM firewall devices. Centralized management not only makes it easier for network managers to configure and modify the security policy, but also makes the firewalling scheme scalable to large MPOA systems.

- It has no impact on performance.

  The proposed firewalling scheme has no performance impact on MPOA system because neither MPOA servers nor edge devices will be slowed down by call screening operations. In MPOA servers, call screening decision making can be implemented efficiently using hashing techniques and strong authentication can be performed very fast using cryptographic hardware [7]. In edge devices, to perform call screening, table lookup is now based on both source address and destination address instead of destination address alone in a normal switch. However, with hashing techniques, the cost difference between them is virtually negligible. Therefore, the proposed firewalling scheme has no impact on performance.

- The proposed scheme is multi-protocol.

  The proposed scheme is especially suitable for MPOA because it is multi-protocol. The firewalling scheme does not assume the internetwork layer protocol being filtered. When

the edge device is about to establish a cut-through connection, it knows which internetwork layer protocol will run over the connection. So it will choose a firewall switch that is able to filter the protocol. The firewalling scheme works not only for TCP/IP, but also for other internetwork layer protocols such as IPX and AppleTalk, as long as ATM firewall switches that are capable of filtering such protocols are available.

4 Conclusion

In this paper, we propose a firewalling scheme for MPOA-based enterprise intranets that is tight in security, low in cost, high in performance, and easy to manage. It is the only secure and cost-effective firewalling approach proposed for MPOA so far. Different from traditional firewalling approaches, it does not assume the existence of physical chokepoints inside a network and places no limitations on network topology, which makes the scheme applicable to all MPOA-based enterprise networks. The firewalling scheme facilitates centralized management, which makes the firewalling scheme scalable to very large MPOA-based systems. The firewalling scheme is "multi-protocol" in the sense that it not only works for TCP/IP, but also works for other internetwork layer protocols as well.

References

[1] A. Alles, "ATM Internetworking", http://www.cisco.com/, Cisco Systems, Inc., May 1995.

[2] E. Andrews, "ATM Migration: Legacy Interoperability Is Key", http://www.lantimes.com/, *LAN Times*, 4/1/96.

[3] D. Chapman, and E. Zwicky, Building Internet Firewalls, O'Reilly & Associates , Sebastopol, CA, 1995.

[4] T. Doty, "The Firewall Heresies", http://www.network.com/, Storage Technology Corp., 1997.

[5] B. Ellington, Ed., "LAN Emulation Over ATM: Version 1.0 Specification", ATM Forum, January 1995.

[6] A. Fredette, Ed., "Multi-Protocol Over ATM Version 2.0", ATM Forum, July, 1997.

[7] IBM Corporation, "IBM 4758 PCI Cryptographic Coprocessor", http://www.ibm.com/, 1998.

[8] J. Jeffords, Ed., "Private Network-Network Interface Specification Version 2.0 (PNNI 2.0)", ATM Forum, September, 1997.

[9] D. Katz, D. Piscitell, B. Cole, and J. Luciani, "NBMA Next Hop Resolution Protocol (NHRP)", Internet Draft (raft-ietf-rolc-nhrp-10.txt), expires March 1997.

[10] B. J. Kowalski, "ATLAS Policy Cache Architecture", http://www.network.com/, Storage Technology Corp., 1997.

[11] D. Minoli, and A. Alles, LAN, ATM, and LAN Emulation Technologies, Artech House, Inc., Norwood, MA, 1996.

[12] D. Newman, and B. Melson, "Can Firewalls Take the Heat?", http://www.data.com/, *Data Communications*, 1998.

[13] D. Passmore, and J. Freeman, "The Virtual LAN Technology Report", http://www.3com.com/, 3com corporation, 1997.

[14] B. Schneier, Applied Cryptography, 2nd ed., John Wiley & Sons, 1996.

[15] J. Xu, and M. Singhal, "Design of A High-Performance ATM Firewall", Ohio State Univeristy Technical Report (OSU-CISRC-4/98-TR13), April, 1998.