A Firewalling Scheme for Securing MPOA-based Enterprise Networks *

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

April 1999

Abstract

A well-known security problem with MPOA is that cut-through connections generally bypasses firewall routers if there are any. None of the previously proposed approaches solved the problem properly. In this paper, we propose a novel firewalling scheme for MPOA that nicely fixes the security hole. Our firewalling scheme has three outstanding advantages that make it ideal for securing MPOA-based enterprise networks. First, based on our novel concept of "logical chokepoints", our firewalling scheme does not require the existence of physical chokepoints inside the network. Second, the scheme is nicely embedded into the MPOA protocol so that its cost, performance overhead, and protocol complexity are reduced to a minimum. Third, the scheme is centrally administrate-red so that it scales well to very large networks.

1 Introduction

MPOA is proposed as a unified framework that allows MAC and internetwork layer protocols be transparently transported over an ATM network[2]. 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 [1], which provide a *virtual LAN* (VLAN) service, where ports on an *edge device* can be assigned to particular VLANs, independent of the physical location. Communication within a VLAN occurs across the ATM network using LANE; communication across VLANs may either travel through the *default path* (an MPOA server or an ATM router), or through *cut-through connections*. A major security hole of MPOA is that cut-through connections generally bypass firewall routers if there are any. Two firewalling schemes for securing cut-through connections have been considered by ATM forum but none of them is

*This work is partially supported by NSA Grant MDA904-96-1-0111. An early version of this paper appeared in Proceedings of IEEE HASE'98.

satisfactory. The first approach is to let each edge device implement packet filtering function[7]. This approach is later refuted by ATM forum due to its high cost. The cost-effectiveness of MPOA lies in the low cost of edge devices. 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 [7], 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 connection only when the first several packets of a *flow* withstand the scrutiny [7]. 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 networks that 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. Unsafe connections will be forced to 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 for two reasons. First, in MPOA, the boundary between security domains (VLANs) are blurred due to the existence of cut-through connections over the ATM network. Second, it is undesirable to have a physical chokepoint inside the ATM network because it will be the bottleneck on the backbone.

The paper is organized as follows. 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-3 (150 Mbps), is a prerequisite of our firewalling scheme. At the time of writing, one commercial ATM firewall product, StorageTek's

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

ATLAS, is available [6]. In [10], we also propose an ATM firewall architecture which nicely embeds high-performance packet-filtering mechanisms into a normal ATM switch. In the following discussion we assume the availability of high-performance ATM firewall devices 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 the 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. The first task is to deploy ATM firewall switches at the "strategic" points of the ATM network of an MPOA system. 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, our firewalling scheme will force it to travel through a firewall switch and get filtered. Such a connection is called a *filtered connection*, as opposite to a *trusted connection*, the content of which needs no inspection. The third task is to establish such a trusted connection when needed. Fig. 1 illustrates our firewalling scheme in a sample MPOA system (oversimplified).

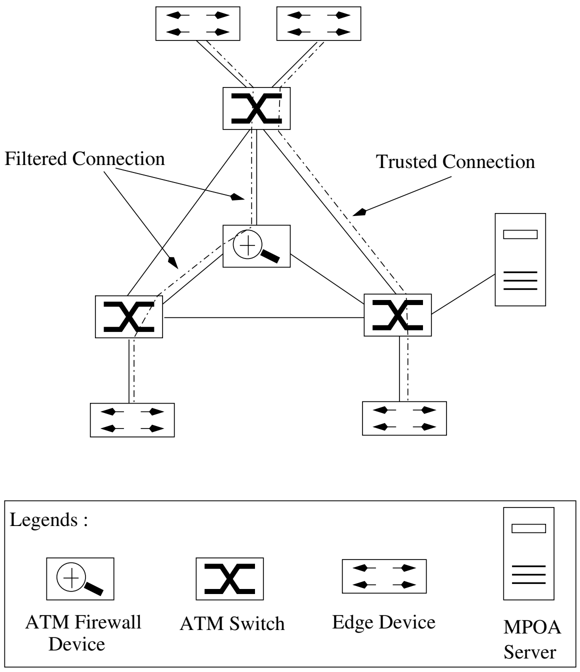

Figure 1: Our firewalling scheme in a sample MPOA system

The basic idea of our firewalling scheme is that edge devices act together to form a protection shield against internal security breaches. 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 [8]. We also 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, the call screening decision is made by MPOA servers. Edge devices interact with MPOA servers to request and obtain call screening decisions. This process is embedded into the MPOA protocol as follows[2]:

1. When edge device A wants to establish a cut-through connection with edge device B, A will send to the MPOA server an 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[3] as well. If *strong authentication* is needed to establish a cut-through connection, authentication information such as a signed digital signature [9] is also contained.

2. When the 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. 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*.

[2]Here, for ease of discussion, we assume that both the ingress edge device and the egress edge device are served by the same MPOA server. In case they are served by different MPOA servers, NHRP flow [5] between MPOA servers can be extended to synchronize information known to both MPOA servers.

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

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 previous section 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* [4]. 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 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:

- It is secure and flexible.

  It is secure because edge devices form a protection shield that forces each and every unsafe connection between different VLANs to 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 to any MPOA-based enterprise networks.

- Its cost 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 also shown in our technical report [11] 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 [10], the cost to implement 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. Therefore, 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 [3]. 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.

4 Conclusion

In this paper, we propose a firewalling scheme for MPOA-based enterprise networks 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.

References

[1] B. Ellington. *LAN Emulation Over ATM: Version 1.0 Specification*. ATM Forum, Jan. 1995.

[2] A. Fredette. *Multi-Protocol Over ATM Version 2.0*. ATM Forum, July 1997.

[3] IBM Corp., http://www.ibm.com/. *IBM 4758 PCI Cryptographic Coprocessor*, 1998.

[4] J. Jeffords. *Private Network-Network Interface Specification Version 2.0 (PNNI 2.0)*. ATM Forum, Sept. 1997.

[5] D. Katz, P. Piscitello, B. Cole, and L. Luciani. NBMA Next Hop Resolution Protocol (NHRP). Internet Draft (raft-ietf-rolc-nhrp-10.txt). expires Mar 1997.

[6] B. Kowalski. ATLAS Policy Cache Architecture. Technical report, StorageTek Corp., http://www.network.com/, 1997.

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

[8] D. Passmore and J. Freeman. The Virtual LAN Technology Report. Technical report, 3com Corp., http://www.3com.com/, 1997.

[9] B. Schneier. *Applied Cryptography*. John Wiley & Sons, NY, 2nd edition, 1996.

[10] J. Xu and M. Singhal. Design of a high-performance atm firewall. In *Proc. of 5th. ACM Conference on Computer and Communication Security*, San Francisco, CA, Nov. 1998.

[11] J. Xu and M. Singhal. A Firewalling Scheme for Securing MPOA-based Corporate Intranets. Technical report, Department of CIS, The Ohio State University, OSU-CISRC-6/98-TR19, June 1998.