Network Working Group
Internet Draft
draft-keromytis-disfi-compare-00.txt

A. D. Keromytis
Columbia University

Some Problems with Perimeter Firewalls

Status of this Memo

Abstract

This document discusses some of the shortcomings of perimeter firewalls and the reasons for employing end-point (or distributed) firewall functionality in the network, either as an alternative or coexisting with traditional network access controls.

1. Introduction

Distributed firewalls [Bel99,IKBS] represent an alternative network access control mechanism to "traditional" perimeter firewalls [BC94]. Policy is enforced in a decentralized manner by the different components of the network, acting in a coordinated manner. Distributed network access control (DNAC) can be employed either as an alternative or in conjunction with perimeter firewalls, and can employ different mechanisms and protocols, as well as heterogeneous elements.

This document discusses some of the shortcomings of perimeter firewalls and the reasons for employing end-point (or distributed) firewall functionality in the network, either as an alternative or coexisting with traditional network access controls.

1.1. Terminology

In this document, the key words "MAY", "MUST, "MUST NOT", "optional", "recommended", "SHOULD", and "SHOULD NOT", are to be interpreted as described in [RFC2119].

2. Shortcomings of Perimeter Firewalls

In this section, we list several shortcomings of Perimeter Firewalls (PFs) and briefly discuss them. It is important to bear in mind that PFs fundamentally depend on restrictions in the network topology, such that they can examine all traffic exchanged between a protected network and the public network, and enforce a security policy (typically access control).

2.1. Performance

Because PFs depend on restrictions in network topology to enforce a security policy, it has to examine all traffic exchanged between the protected network and the public network. Thus it can become a performance bottleneck, especially if it also has to perform network encryption to accomodate Virtual Private Networks (VPNs), virus checking, or other CPU-intensive services. Load-balancing techniques can help to some extend, but the complexity of state-sharing and synchronization in firewall clusters is a real limiting factor.

2.2. Protection from Insiders

Likewise, because of the dependence on the network topology, a PF can only enforce a policy on traffic that traverses it. Thus, traffic exchanged among nodes in the protected network cannot be controlled. This gives an attacker that is already an insider or can somehow bypass the firewall (see Section 2.6) complete freedom to act. One obvious solution is to deploy multiple firewalls, in which case they must be coordinated and control in a consistent manner.

2.3. End-to-end Protocol Properties

There are several stateful protocols that use random ports for data transfers, which require the PF to actively monitor (and reconstruct) the control message sequences for these protocols. This increases the complexity and decreases the performance of the PF, and thus the whole (protected) network. The fundamental problem is that IP was designed as a lean network substrate that would simply act as a packet carrier; all protocol state is maintained at the communication end-points. A PF breaks this model by definition (since it needs to peek inside packets to make a policy decision); for some protocols, the ability to make that decision requires reconstruction of part of the state that is

readily available at the end-points.

2.4. Network Encryption

Increasingly, end-to-end cryptographic protocols such as IPsec [RFC2401], TLS [RFC2246], and SSH are being used to protect (encrypt) communications between nodes. PFs are hard-pressed to enforce a security or intrusion detection (e.g., virus scanning) policy in these situations, as they cannot strip the encryption. Potential solutions break this end-to-end model of security (key escrow, firewall as encryption end-point or trusted man-in-the-middle, etc.) with a corresponding loss of assurance and performance.

2.5. Mobility / Telecommuting / Infrastructure Sharing

In certain scenarios, nodes or users that are topologically outside the protected network must be treated as if they were inside it. Examples include users on the road, telecommuters, and Intranet configurations. In all these cases, parts (or all) of the protected infrastructure need to be exposed in a controlled manner to remote entities. PFs make this a complicated subject, especially in the presence of multihoming and fine-grain resource sharing (and corresponding access control) requirements.

2.6. Closure

Due to several trends in networking such as site multihoming, abundant dial-up access, and wide-scale (insecure) wireless deployment, administrators cannot be assured that all traffic exchanged between the protected and the public network will be examined by a PF: an attacker that manages to access the protected network through a user-installed wireless access point has complete freedom of action.

2.7. Defense in Depth

Following classic army doctrine, prudent security engineering argues for multiple defensive lines ("defense in depth") as a way to minimize the risk of a successful breach of the perimeter. PFs alone cannot (by definition) provide this functionality; additional mechanisms are necessary inside the protected network.

3. Security Considerations

This draft describes some shortcomings of perimeter firewalls, and argues for the need for distributed firewall functionality. There are several security issues in such an architecture, including policy robustness, secure distribution and control, coordination, management, and resilience to intrusions or insider attacks (this is not intended to be a detailed or exclusive list). These issues

will be addressed in separate documents of the working group (should there be one).

4. IANA Considerations

No requirements are placed on IANA at this stage.

Acknowledgements

This document came out of discussions at the Distributed Firewalls mailing list, which can be foiund at disfi@lists.intel.com

To subscribe on the list, send email to listserv@lists.intel.com with the command "subscribe disfi" in the message body.

References:

[Bel99] S.M. Bellovin, "Distributed Firewalls", USENIX ;login magazine, special issue on security, November 1999.

[BC94] Bellovin, S.M. and W.R. Cheswick, "Firewalls and Internet Security: Repelling the Wily Hacker", Addison-Wesley, 1994.

[IKBS] Ioannidis, S. and Keromytis, A.D., and Bellovin, S.M. and J.M. Smith, "Implementing a Distributed Firewall", Proceedings of Computer and Communications Security (CCS), pp. 190-199, November 2000, Athens, Greece.

[RFC 2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2246] T. Dierks and C. Allen, "The TLS Protocol Version 1.0," Internet Engineering Task Force, RFC 2246, January 1999.

[RFC2401] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.

Author's address:

Angelos D. Keromytis
Columbia University, CS Department
515 CS Building
1214 Amsterdam Avenue, Mailstop 0401
New York, New York 10027-7003

Phone: +1 212 939 7095
Email: angelos@cs.columbia.edu

Expiration and File Name

This draft expires in May 2003

Its file name is draft-keromytis-disfi-compare-00.txt