Chapter 1

Introduction

A challenging problem in IP Network is how we can protect our network infrastructure from malicious traffic, such as scanning, malicious code propagation, spam, and distributed denial-of-service (DDoS) attacks. These activities cause problems on a regular basis, ranging from simple annoyance to severe financial, operational and political damage to companies, organizations, and critical infrastructure. In recent years, they have increased in volume, sophistication, and automation, largely enabled by botnets, which are used as the platform for launching these attacks [11].

Protecting a victim (host or network) from malicious traffic is a hard problem that requires the coordination of several complementary components, including nontechnical (e.g., business and legal) and technical solutions (at the application and/or network level). Filtering support from the network is a fundamental building block in this effort. For example, an Internet service provider (ISP) may use filtering in response to an ongoing DDoS attack to block the DDoS traffic before it reaches its clients. Another ISP may want to proactively identify and block traffic carrying malicious code before it reaches and compromises vulnerable hosts in the first place. In either case, filtering is a necessary operation that must be performed within the network. Filtering capabilities are already available at routers today via access control lists (ACLs). ACLs enable a router to match a packet header against predefined rules and take predefined actions on the matching packets, and they are currently used for enforcing a variety of policies, including infrastructure protection.

For the purpose of blocking malicious traffic, a filter is a simple ACL rule that denies access to a source IP address or prefix. To keep up with the high forwarding rates of modern routers, filtering is implemented in hardware: ACLs are typically stored in ternary content addressable memory (TCAM), which allows for parallel access and reduces the number of lookups per forwarded packet. However, TCAM is more expensive and consumes more space and power than conventional memory. The size and cost of TCAM puts a limit on the number of filters, and this is not expected to change in the near future. With thousands or tens of thousands of filters per path, an ISP alone cannot hope

to block the currently witnessed attacks, not to mention attacks from multimillion-node botnets expected in the near future.

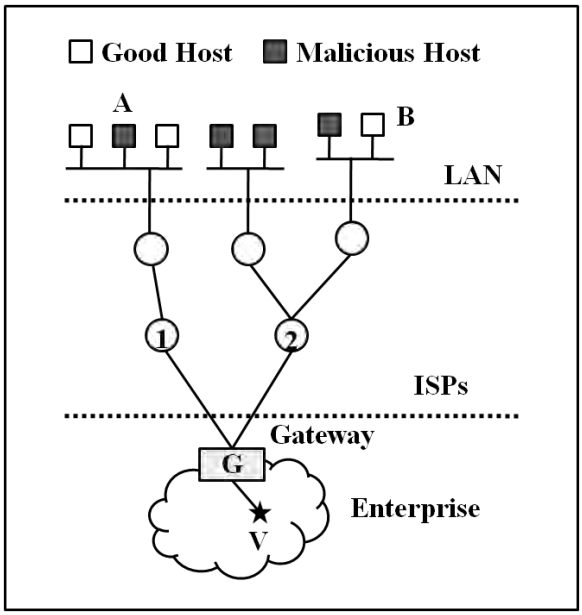

(a)

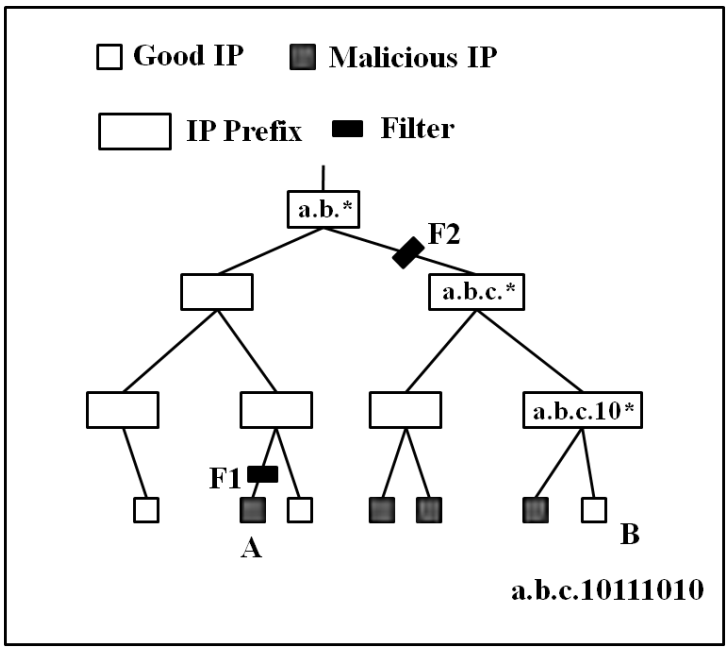

(b)

Figure 1.1: Example of a distributed attack. (a) Actual Network and (b) Hierarchy of Source IP Addresses and Prefixes [11].

Consider the example as shown in Figure 1.1(a): an attacker commands a large number of compromised hosts to send traffic towards a victim V (say a web server), thus exhausting the resources of V and preventing it from serving its legitimate clients; the ISP

of V tries to protect its client from the attack, by blocking the malicious traffic at the gateway router G. Ideally, G would like to assign a single filter to block each malicious IP source. However, there are fewer filters than attackers so aggregation is typically used: a single filter blocks an entire source address prefix. This has the desired effect of reducing the number of filters, but also the undesired side-effect of blocking legitimate traffic originating from that prefix.

In figure 1.1(b): Let us assume that the gateway router G has only two filters available to block malicious traffic and protect the victim V . It uses F1 to block a single malicious address (A) and F2 to block prefix a.b.c.*, which contains 3 malicious sources but also one legitimate source (B). Therefore, the selection of filter F2 trades-off the collateral damage (blocking B) for the reduction in the number of filters.

1.1 Background

In this section we will discuss about the firewalls, the need for firewalls, firewall characteristics and types of firewalls.

1.1.1 Firewall

A firewall can either be software-based or hardware-based and is used to help keep a network secure. Its primary objective is to control the incoming and outgoing network traffic by analyzing the data packets and determining whether it should be allowed through or not, based on a predetermined rule set. A network's firewall builds a bridge between an internal network that is assumed to be secure and trusted, and another network, usually an external (inter) network, such as the Internet, that is not assumed to be secure and trusted (see in figure 1.2).

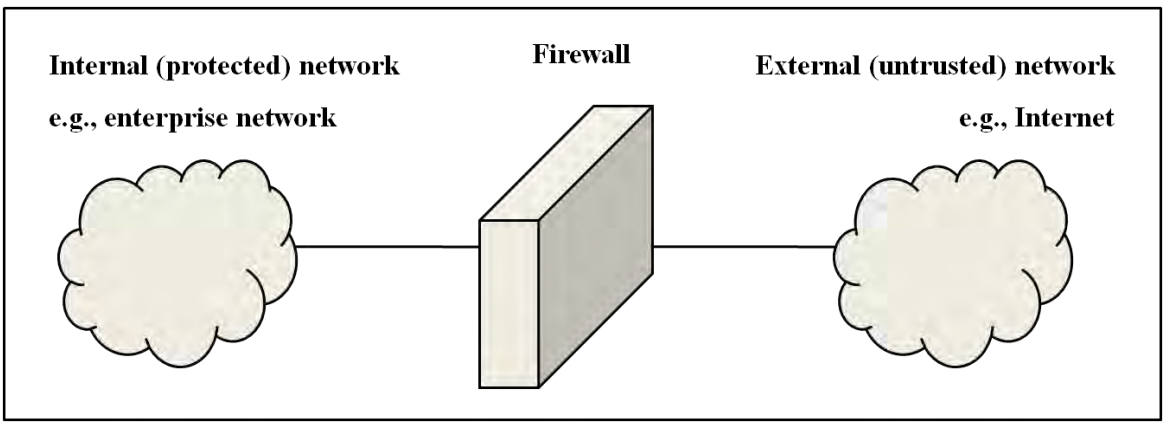

Figure 1.2: General Model of Firewall [12].