CMSGu2014-02

National Computer Board

Mauritian Computer Emergency Response Team

Enhancing Cyber Security in Mauritius

Guideline on Firewall

CERT-MU

National Computer Board
Mauritius

June 2014

Issue No. 2

Version 1.0

Table of Contents

DISCLAIMER: *This guideline is provided "as is" for informational purposes only. Information in this guideline, including references, is subject to change without notice. The products mentioned herein are the trademarks of their respective owners.*

1.0 Introduction

1.1 Purpose and Scope

The purpose of this guideline is to give organisations an insight of the different firewall technologies that are available and the different firewall architectures that could be applied to protect the network of an organisation.

1.2 Audience

The target audience for this document includes CIOs, CISOs, information security staffs, network administrators and all other relevant parties involved in the maintenance of the IT infrastructure.

1.3 Document Structure

This document is organised into the following sections:
Section 1 provides an overview on the document's content, the targeted audience and the document's structure.

Section 2 gives background on firewall.

Section 3 presents an overview of firewall technologies.

Section 4 gives a description of the most common firewall and network architectures that an organisation might use.

Section 5 gives an insight of some of the most important firewalls rules.

Section 6 concludes the document.

Section 7 consists of a list of references that have been used in this document.

Appendix A provides a list of acronyms that have been used in the document.

2.0 Background

Firewalls are devices that allow or block traffic into and out of a private network or the user's computer. Firewalls are frequently used to prevent unauthorised Internet users from accessing private networks connected to the Internet. All data entering or leaving the Intranet pass through the firewall, which examines each packet and blocks those that do not meet the specified security criteria.

Generally, firewalls are configured to protect against unauthenticated interactive logins from the outside world. This helps prevent hackers from logging into machines on your network. Firewalls are essential since they can provide a single block point where security and auditing can be imposed. Firewalls also provide an important logging and auditing function; often they provide summaries to the administrator about what type/volume of traffic has been processed through it.

In an organisation, a firewall can be a stand-alone machine or software in a router or server. It can be as simple as a single router that filters out unwanted packets, or it may comprise a combination of routers and servers each performing some type of firewall processing.

In the home, a personal firewall typically comes with or is installed in the user's computer. Personal firewalls may also detect outbound traffic to guard against spyware, which could be sending your surfing habits to a Web site. They may also alert you when software makes an outbound request for the first time.

3.0 Overview of Firewall Technologies

There are several types of firewall technologies that are available and one way of comparing their capabilities is to look at the Transmission Control Protocol/Internet Protocol (TCP/IP) layers that each is able to examine. Firewalls are often placed at the perimeter of a network and such a firewall is said to have an external and internal interface, with the external interface being the one on the outside of the network. These two interfaces are also referred to as unprotected and protected, respectively.

This section provides an overview of firewall technologies and basic information on the capabilities of several commonly used types.

3.1 Packet Filtering Firewalls

This technology belongs to the first generation of firewalls. It works at network and transport layer in the OSI model (layer 3 and 4), analysing IP addresses and ports. Each packet that enters or leaves the network is inspected and accepted or rejected based on the rules defined by the firewall administrator. Packet filtering is effective and transparent to network users. Decisions based on packet filtering are taken rapidly, and therefore this type of firewall offers optimum performance. However, there are some problems associated with this type of firewall for e.g. (1) they cannot determine if the packet that it has let through contains some type of malicious code, (2) they are difficult to setup and configure and (3) they are vulnerable to IP Spoofing.

3.2 Stateful Packet Filtering Firewalls

This technology is from the second generation of firewalls and validates that packets correspond to a connection request or to a connection between two devices. It applies security mechanisms when a TCP or UDP connection is established.

Stateful packet filters keep an internal table with the state of the connections through the firewall. This type of firewall decides whether to accept or reject traffic on a connection-by-connection basis. These decisions are taken using both the information used by simple packet filters and the internal connections filter. The performance is also optimum, even better than that of the simple packet filter, as in order to decide what to do with a packet corresponding to a connection; it only needs to consult the table. Once the connection has ended, its entry is deleted from the state table and data transmission is closed.

However, the connections table requires memory space, and they should therefore be run on systems with adequate memory space.

3.3 Application Firewalls

A newer trend in stateful inspection is the addition of a stateful protocol analysis capability, referred to by some vendors as deep packet inspection. Stateful protocol analysis improves upon standard stateful inspection by adding basic intrusion detection technology—an inspection engine that analyses protocols at the application layer to compare vendor-developed profiles of benign protocol activity against observed events to identify deviations. This allows a firewall to allow or deny access based on how an application is running over the network. For instance, an application firewall can determine if an email message contains a type of attachment that the organisation does not permit (such as an executable file), or if instant messaging (IM) is being used over port 80 (typically used for HTTP). Another feature is that it can block connections over which specific actions are being performed (e.g., users could be prevented from using the FTP "put" command, which allows users to write files to the FTP server). This feature can also be used to allow or deny web pages that contain particular types of active content, such as Java or ActiveX, or that have SSL certificates signed by a particular certificate authority (CA), such as a compromised or revoked CA.

3.4 Application-Proxy Gateways

An application-proxy gateway is a feature of advanced firewalls that combines lower-layer access control with upper-layer functionality. These firewalls contain a proxy agent that acts as an intermediary between two hosts that wish to communicate with each other, and never allows a direct connection between them. Each successful connection attempt actually results in the creation of two separate connections—one between the client and the proxy server, and another between the proxy server and the true destination. The proxy is meant to be transparent to the two hosts—from their perspectives there is a direct connection. Because external hosts only communicate with the proxy agent, internal IP addresses are not visible to the outside world. The proxy agent interfaces directly with the firewall ruleset to determine whether a given instance of network traffic should be allowed to transit the firewall.

Application-proxy gateways are quite different than application firewalls. First, an application-proxy gateway can offer a higher level of security for some applications because it prevents direct connections between two hosts and it inspects traffic content to identify policy violations. Another potential advantage is that some application-proxy gateways have

the ability to decrypt packets (e.g., SSL-protected payloads), examine them, and re-encrypt them before sending them on to the destination host. Data that the gateway cannot decrypt is passed directly through to the application. When choosing the type of firewall to deploy, it is important to decide whether the firewall actually needs to act as an application proxy so that it can match the specific policies needed by the organisation.

3.5 Unified Threat Management (UTM)

Many firewalls combine multiple features into a single system, the idea being that it is easier to set and maintain policy on a single system than on many systems that are deployed at the same location on a network. A typical unified threat management (UTM) system has a firewall, malware detection and eradication, sensing and blocking of suspicious network probes, and so on. There are pros and cons to merging multiple, not-completely-related functions into a single system. For example, deploying a UTM reduces complexity by making a single system responsible for multiple security objectives, but it also requires that the UTM have all the desired features to meet every one of the objectives. Another tradeoff is in performance: a single system handling multiple tasks has to have enough resources such as CPU speed and memory to handle every task assigned to it. Some organisations will find the balance favors a UTM, while other organisations will use multiple firewalls at the same location in their network.

3.6 Web Application Firewalls

The HTTP protocol used in web servers has been exploited by attackers in many ways, such as to place malicious software on the computer of someone browsing the web, or to fool a person into revealing private information that they might not have otherwise. Many of these exploits can be detected by specialised application firewalls called web application firewalls that reside in front of the web server.

Web application firewalls are a relatively new technology, as compared to other firewall technologies, and the type of threats that they mitigate are still changing frequently. Because they are put in front of web servers to prevent attacks on the server, they are often considered to be very different than traditional firewalls.

4.0 Firewall and Network Architectures

The main use of firewalls is to separate networks that have different security requirements, such as the Internet and an internal network of an organisation that is hosting sensitive information. Organisations should place a firewall whenever its internal networks and systems interface with external networks and systems, and where security requirements vary among their internal networks. This section is intended to help organisations determine where firewalls should be placed, and where other networks and systems should be located in relation to the firewalls.

4.1 Dual-Homed Firewall

A dual-homed firewall uses two (or more) network interfaces. One connection is an internal network (trusted network) and the second connection is to the Internet (untrusted network). The key security principle is not to allow traffic coming in from an untrusted network to be directly routed to the trusted network. In this scenario, the IP source routing and IP forwarding services are disabled.

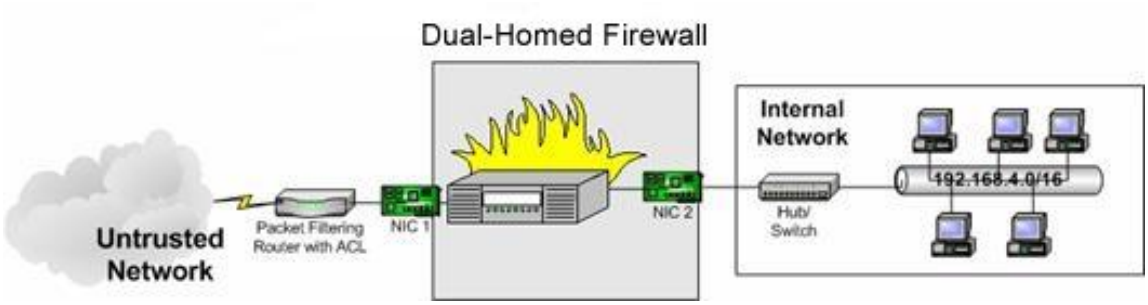

Figure 1: Dual-Homed Firewall

4.2 Screened Host Architecture

In the screened host architecture, the bastion host sits on the internal network. The packet filtering on the screening router is set up in such a way that the bastion host is the only system on the internal network that hosts on the Internet can open connections to (for example, to deliver incoming email). Even then, only certain types of connections are allowed. Any external system trying to access internal systems or services will have to connect to this host. The bastion host thus needs to maintain a high level of host security.

However, compared to other architectures such as the DMZ architecture which will be discussed in the following section, there are some disadvantages associated to the screen host architecture. The major one is that if an attacker manages to break in to the bastion host,

there is nothing left in the way of network security between the bastion host and the rest of the internal hosts. The router also presents a single point of failure; if the router is compromised, the entire network is available to an attacker. For this reason, the screened host architecture has not been widely implemented.

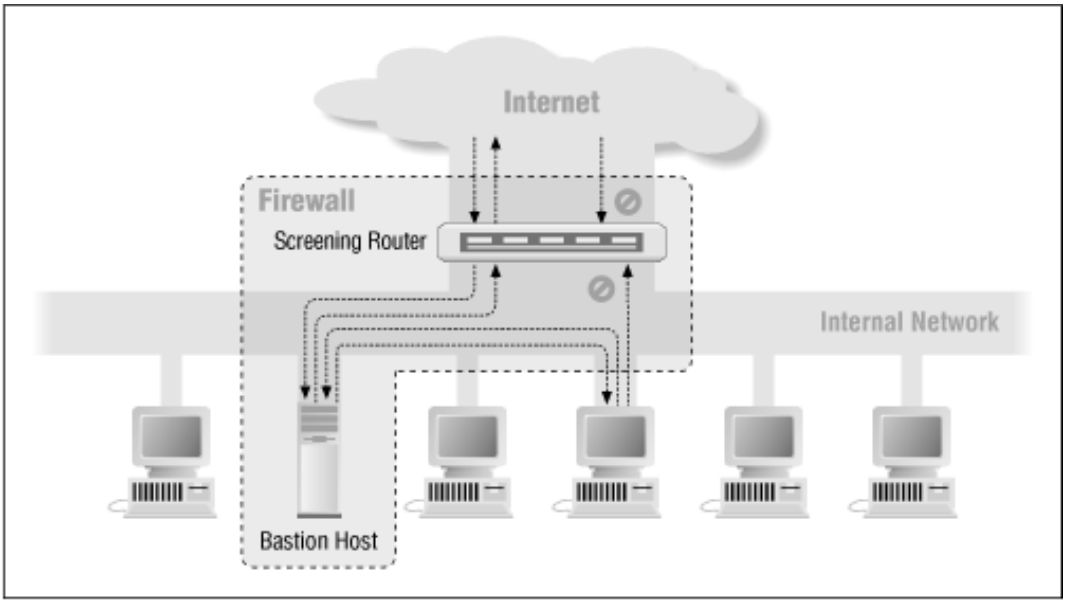

Figure 2: Screened Host Architecture

4.2 DMZ Architecture

A demilitarised zone (DMZ) architecture isolates hosts which are accessible from outside the network (e.g. a web server or FTP server) from internal servers. The external hosts are placed in a separate network zone, on a separate adapter, connected to the firewall. This is easily achieved with a firewall with three or more interfaces or by using two distinct firewalls.

Each subnet is also configured with its own security zone (e.g. the Finance network, the Sales network, etc.) by connecting it to a separate firewall adapter. All traffic between zones, and all traffic from the Internet to all zones, is checked by the firewall.

In this way, each zone is isolated, and the systems in each zone only trust other systems within the same zone. Therefore, if a hacker succeeds in breeching an accessible host, the other hosts within the network are still safe.

DMZs are often used for special servers, such as web servers, which must be accessible from two separate networks. Usually an organisation has one Internet connection, one local network and one DMZ with servers that must be both internally and externally accessible. Figure 3 shows a typical DMZ architecture.

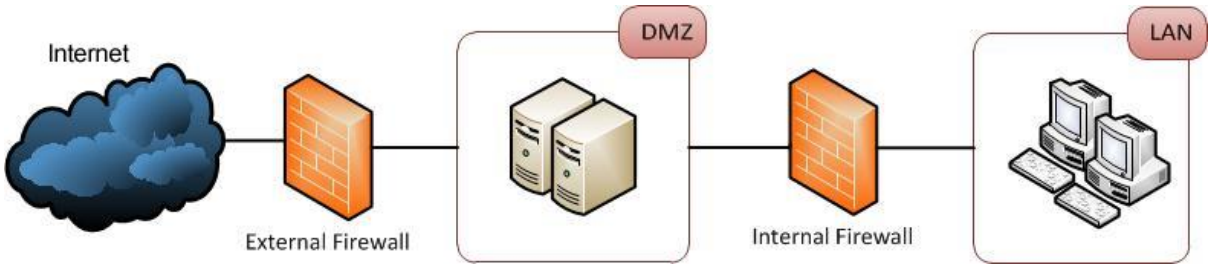

Figure 3: DMZ Architecture

5.0 Firewall Policy

A firewall policy is a set of rules on how a firewall should handle network traffic for specific IP addresses, address ranges, protocols, applications and content types based on the organisation's information security policies. Before creating a firewall policy, some form of risk analysis should be performed to put out a list of the types of traffic needed by the organisation and categorise on how they must be secured. This section provides details on what types of traffic should be blocked.

4.2 Firewall Filtering Rules

Firewall policies should only permit appropriate source and destination IP addresses to be used. Below shows some specific recommendations for IP addresses:

- Traffic with invalid source or destination addresses should always be blocked, regardless of the firewall location. Examples of relatively common invalid IPv4 addresses are 127.0.0.0 to 127.255.255.255 (also known as the localhost addresses) and 0.0.0.0 (interpreted by some operating systems as a localhost or a broadcast address). Also, traffic using link-local addresses (169.254.0.0 to 169.254.255.255) should be blocked.
- Traffic with an invalid source address for incoming traffic or destination address for outgoing traffic (an invalid "external" address) should be blocked at the network perimeter. This traffic is often caused by malware, spoofing, denial of service attacks, or misconfigured equipment. The most common type of invalid external addresses is an IPv4 address within the ranges in RFC 1918, *Address Allocation for private Internets*, which are reserved for private networks. These ranges are 10.0.0.0 to 10.255.255.255 (10.0.0.0/8 in Classless Inter-Domain Routing [CIDR] notation), 172.16.0.0 to 172.31.255.255 (172.16.0.0/12), and 192.168.0.0 to 192.168.255.255 (192.168.0.0/16).
- Traffic with a private destination address for incoming traffic or source address for outgoing traffic (an "internal" address) should be blocked at the network perimeter. Perimeter devices can perform address translation services to permit internal hosts with private addresses to communicate through the perimeter, but private addresses should not be passed through the network perimeter.
- Outbound traffic with invalid source addresses should be blocked (this is often called *egress filtering*). Systems that have been compromised by attackers can be used to

attack other systems on the Internet; using invalid source addresses makes these kinds of attacks more difficult to stop. Blocking this type of traffic at an organisation's firewall helps reduce the effectiveness of these attacks.

- Incoming traffic with a destination address of the firewall itself should be blocked unless the firewall is offering services for incoming traffic that require direct connections—for example, if the firewall is acting as an application proxy.
- Traffic containing IP source routing information, which allows a system to specify the routes that packets will employ while traveling from source to destination should be blocked. This could potentially permit an attacker to construct a packet that bypasses network security controls. IP source routing is rarely used on modern networks, and valid applications are even less common on the Internet.
- Traffic from outside the network containing broadcast addresses that are directed inside the network should also be blocked. Any system that responds to the directed broadcast will then send its response to the system specified by the source, rather than to the source system itself. These packets can be used to create huge "storms" of network traffic for denial of service attacks. Regular broadcast addresses, as well as addresses used for multicast IP, may or may not be appropriate for blocking at an organisation's firewall. Multicast and broadcast networking is seldom used in normal networking environments, but when it is used both inside and outside of the organisation, it should be allowed through firewalls.
- Firewalls at the network perimeter should block all incoming traffic to networks and hosts that should not be accessible from external networks and should also block all outgoing traffic from the organisation's networks and hosts that should not be permitted to access external networks.

6.0 Conclusion

Firewalls have been around for a certain number of years and it has become an integral part of the network for providing network security as it is usually the first layer of security controls in place to protect an organisation's network. Therefore, choosing the right type of firewall together with the proper implementation and the setting up a strong firewall policy are crucial factors for the successful deployment of a firewall in an organisation. Nowadays, there are a number of Enterprise Network Firewalls that are available on the market and the common ones as stated by the Gartner Magic Quadrant are Check Point Software Technologies, Palo Alto Networks, Fortinet, Juniper and Cisco.

7.0 References

- NIST, Guidelines on Firewall and Firewall Policy, **csrc.nist.gov**
- TechTarget, Introduction to Firewalls: Types of Firewall, **techtarget.com**
- Next-Generation Firewalls for Dummies by Lawrence C. Miller
- Firewalls: A Technical Overview, **boran.com**

Appendix A

List of Acronyms

CIDR	Classless Inter-Domain Routing
CIO	Chief Information Officer
CISO	Chief Information Security Officer
DMZ	Demilitarised Zone
IP	Internet Protocol
OSI	Open System Interconnection
RFC	Request for Comments
TCP	Transmission Control Protocol
UTM	Unified Threat Management