

Operations Area Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: August 7, 2016

F. Gont
SI6 Networks / UTN-FRH
F. Baker
Cisco Systems
February 4, 2016

On Firewalls in Network Security
draft-gont-opsawg-firewalls-analysis-02

Abstract

This document analyzes the role of firewalls in network security, and recognizes their role in the internet architecture. It suggests a line of reasoning about their usage, and analyzes common kinds of firewalls and the claims made for them.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 7, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Terminology	3
3. Reasoning about Firewalls	4
3.1. A Simple Model of Communication	4
3.2. The Role of Firewalls in Internet Security	5
3.3. Firewalls and The End-to-End Principle	5
4. Common kinds of firewalls	6
4.1. Perimeter security: Protection from aliens and intruders	7
4.2. Pervasive access control	8
4.3. Intrusion Management: Contract and Reputation filters . .	9
5. Firewalling Strategies	10
5.1. Blocking Traffic Unless It Is Explicitly Allowed (default deny)	11
5.2. Allow Traffic Unless It Is Explicitly Blocked (default allow)	11
6. Assumptions on IP addresses and Transport Protocol Port Numbers	12
7. State Associated with Filtering Rules	13
8. Enforcing Protocol Syntax at the Firewall	14
9. Performing Deep Packet Inspection	14
10. IANA Considerations	15
11. Security Considerations	15
12. Acknowledgements	15
13. References	16
13.1. Normative References	16
13.2. Informative References	16
Authors' Addresses	18

1. Introduction

Prophylactic perimeter security in the form of firewalls, and the proper use of them, have been a fractious sub-topic in the area of internet security. Firewalls have been largely seen by many in the IETF as a poor approach to security, and often as unnecessary and rather "evil" devices that hinder innovation and the deployment of new protocols and applications. Operationally, they are also seen by some as attack vectors, with state exhaustion attacks, side-effects of the imposition of symmetry requirements and single points of failure. This document analyzes the role of firewalls in network security, and recognizes their role in the internet architecture. It suggests a line of reasoning about their usage, and analyzes common kinds of firewalls and the claims made for them.

This document has, among others, the following goals:

- o Recognize the important role of firewalls in enterprise security architecture for providing "prophylactic" security, rather than as "evil" ad-hoc functionality/devices (see Section 3.2).
- o Analyze common kinds of firewalls and claims made for them (see Section 4).
- o Analyze implicit assumptions made by firewalls, identifying where/when some of those assumptions may not apply (see e.g. Section 6).
- o Discuss trade-offs in the possible firewalling paradigms (see Section 5).
- o Provide conceptual guidance regarding the use and deployment of .
- o Identify harmful behavior/policies commonly implemented and applied by firewalls, in the hopes of improving the state of affairs in that area.
- o Possibly trigger other work in the area of firewalls, as a result of the previous items.

2. Terminology

Firewall:

A device or software that imposes a policy whose effect is "a stated type of network traffic may or may not be allowed from A to B". The firewall may reside in the destination itself (a "host firewall"), or in any intermediate system (a "network firewall"). The firewalling functionality may be implemented in a general purpose system (e.g. an ACL in a router), or in a special purpose middleware device (e.g., a "firewall product"). The details of the policy, the granularity with which a policy can be applied, how such policy is configured, or of the firewall's implementation are just that - implementation details.

We also note that a firewall may enforce policies at different layers. Typically, the layer at which a firewall operates will impact the type of policies that a firewall will be able to apply: for example, a layer-3 firewall may be able to enforce simple policies based on layer-3 addresses and some simple layer-4 parameters such as transport protocol port numbers, while an "application firewall" may be able to enforce policies on higher-level entities such as application-request types. We note that all such firewall types essentially enforce the same role of enforcing a policy of some sort on network traffic, and hence are

referred to with the generic term "firewall" (or "firewall device" in some cases) throughout this document.

Perimeter:

The position in which the specific security policy applies. In typical deployed networks, there are usually some easy- to-define perimeters. A network connected with another network has a perimeter where the two meet, which is defined by what equipment is operated by each network. It invariably imposes a security policy at that boundary, which may be as simple as "all traffic is welcome" and as complex as matching arriving and departing traffic to ensure specific behaviors, or inspecting traffic according to various algorithms. Firewall functionality is usually implemented at or close to such network perimeters.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

3. Reasoning about Firewalls

3.1. A Simple Model of Communication

Any communication requires at least three components:

- o a sender, someone or some thing that sends a message,
- o a receiver, someone or some thing that receives the message, and
- o a channel, which is a medium by which the message is communicated.

In the Internet, the IP network is the channel; it may traverse something as simple as a directly connected cable or as complex as a sequence of ISPs, but it is the means of communication. In normal communications, a sender sends a message via the channel to the receiver, who is willing to receive and operate on it. In contrast, attacks are a form of harassment. A receiver exists, but is unwilling to receive the message, has no application to operate on it, or is by policy unwilling to. Attacks on infrastructure occur when message volume overwhelms infrastructure or uses infrastructure but has no obvious receiver.

By that line of reasoning, a firewall operating at layer-3 primarily protects infrastructure, by preventing traffic that would attack it from it. The best prophylactic might use a procedure for the dissemination of Flow Specification Rules [RFC5575] to drop traffic sent by an unauthorized or inappropriate sender or which has no host

or application willing to receive it as close as possible to the sender.

In other words, a firewall is comparable to the human skin, and has as its primary purpose the prophylactic defense of a network. By extension, the firewall also protects a set of hosts and applications, and the bandwidth that serves them, as part of a strategy of defense in depth. Since there is no one way to prevent attacks, a firewall is not itself a security strategy; the analogy to the skin would say that a body protected only by the skin has an immune system deficiency and cannot be expected to long survive. That said, every security solution has a set of vulnerabilities; the vulnerabilities of a layered defense is the intersection of the vulnerabilities of the various layers (e.g., a successful attack has to thread each layer of defense).

3.2. The Role of Firewalls in Internet Security

One could compare the role of firewalls in prophylactic perimeter security to that of the human skin: the service that the skin performs for the rest of the body is to keep common crud out, and as a result prevent much damage and infection that could otherwise occur. The body supplies prophylactic perimeter security for itself and then presumes that the security perimeter has been breached; real defenses against attacks on the body include powerful systems that detect changes (anomalies) counterproductive to human health, and recognizable attack syndromes such as common or recently-seen diseases. One might well ask, in view of those superior defenses, whether there is any value in the skin at all; the value is easily stated, however. It is not in preventing the need for the stronger solutions, but in making their expensive invocation less needful and more focused.

3.3. Firewalls and The End-to-End Principle

One common complaint about firewalls in general is that they violate the End-to-End Principle [Saltzer]. The End-to-End Principle is often incorrectly stated as requiring that "application specific functions ought to reside in the end nodes of a network rather than in intermediary nodes, provided they can be implemented 'completely and correctly' in the end nodes" or that "there should be no state in the network." What it actually says is heavily nuanced, and is a line of reasoning applicable when considering any two communication layers.

[Saltzer] "presents a design principle that helps guide placement of functions among the modules of a distributed computer system. The principle, called the end-to-end argument, suggests that

functions placed at low levels of a system may be redundant or of little value when compared with the cost of providing them at that low level."

In other words, the End-to-End Argument is not a prohibition against lower layer retries of transmissions, which can be important in certain LAN technologies, nor of the maintenance of state, nor of consistent policies imposed for security reasons. It is, however, a plea for simplicity. Any behavior of a lower communication layer, whether found in the same system as the higher layer (and especially application) functionality or in a different one, that from the perspective of a higher layer introduces inconsistency, complexity, or coupling, extracts a cost. That cost may be in user satisfaction, difficulty of management or fault diagnosis, difficulty of future innovation, reduced performance, or something else. Such costs need to be clearly and honestly weighed against the benefits expected, and used only if the benefit outweighs the cost.

From that perspective, introduction of a policy that prevents communication under an understood set of circumstances, whether it is to prevent access to pornographic sites or to prevent traffic that can be characterized as an attack, does not fail the End-to-End Argument; there are any number of possible sites on the network that are inaccessible at any given time, and the presence of such a policy is easily explained and understood.

What does fail the End-to-End Argument is behavior that is intermittent, difficult to explain, or unpredictable. If a site can be reached sometimes and not at other times, or can be reached using this host or application but not another, one will wonder why that is the case, and may not even know where to look for the issue.

4. Common kinds of firewalls

There are at least three common kinds of firewalls:

- o Context or Zone-based firewalls, that protect systems within a perimeter from systems outside it,
- o Pervasive routing-based measures, which protect intermingled systems from each other by enforcing role-based policies, and
- o Systems that analyze network traffic behavior and trigger on events that are unusual, match a signature, or involve an untrusted peer.

Each kind of firewall addresses a different view of the network. A zone-based firewall (Section 4.1) views the network as containing

zones of trust, and deems applications inside its zone of protection to be trustworthy. A role-based firewall (Section 4.2) identifies parties on the basis of membership in groups, and prevents unauthorized communication between groups. A reputation, anomaly, or signature-based intrusion management system (Section 4.3) depends on active administration, and permits known applications to communicate while excluding unknown or known-evil applications. In each case, the host or application is its own final bastion of defense, but having a host blocking incoming traffic (so-called "host firewalls") does not defend infrastructure. That is, each type of prophylactic has a purpose, and none of them is a complete prophylactic defense.

Each type of defense, however, can be assisted by enabling an application running in a host to inform the network of what it is willing to receive. As noted in Section 4.1, a zone-based firewall, generally denies all incoming sessions and permits responses to sessions initiated outbound from the zone, but can in some cases be configured to also permit specific classes of incoming session requests, such as WWW or SMTP to an appropriate server. A simple way to enable a zone-based firewall to prevent attacks on infrastructure (traffic to an un-instantiated address or to an application that is off) while not impeding traffic that has a willing host and application would be for the application to inform the firewall of that willingness to receive incoming sessions. The Port Control Protocol [RFC6887], or PCP, is an example of a protocol designed for that purpose.

4.1. Perimeter security: Protection from aliens and intruders

As discussed in [RFC6092], the most common kind of firewall is used at the perimeter of a network. Perimeter security assumes two things: that applications and equipment inside the perimeter are under the control of the local administration and are therefore probably doing reasonable things, and that applications and equipment outside the perimeter are unknown.

For example, it may enforce simple permission rules, such as that external web clients are permitted to access a specific web server or that external SMTP MTAs are permitted to access internal SMTP MTAs. Apart from those rules, a session may be initiated from inside the perimeter, and responses from outside will be allowed through the firewall, but sessions may never be initiated from outside.

In addition, perimeter firewalls often perform some level of inspection/analysis, either as application proxies or through deep packet inspection, to verify that the protocol claimed to be being passed is in fact the protocol being passed.

In many scenarios the existence and definition of zone-based perimeter defenses is arguably a side-effect of the deployment of Network Address Translation [RFC2993]. Since e.g. a single address is shared among multiple systems, the NAT device needs to translate both the IP addresses and the transport protocol ports in order to multiplex multiple communication instances from different nodes into the same external address. Thus, the NAT device must keep a state table to know how to translate the IP addresses and transport protocol ports of incoming packets. Packets originating from the internal network will either match an existing entry in the state table, or create a new one. On the other hand, packets originating in the external network will either match an existing entry in the state table, or be dropped. Thus, as a side effect, NATs implicitly require that communication be initiated from the internal network, and only allow return traffic from the external network. We note that this is a side-effect of multiplexing traffic from multiple nodes on a single IP address, rather than a design goal of NAT devices or their associated network translation function.

Some applications make the mistake of coupling application identities to network layer addresses, and hence employ such addresses in the application protocol. Thus, Network Address Translation forces the translator to interpret packet payloads and change addresses where used by applications.

As a result, if the transport or application headers are not understood by the translator, this has the effect of damaging or preventing communication. Detection of such issues can be sold as a security feature, although it is really a side-effect of a failure. While this can have useful side-effects, such as preventing the passage of attack traffic that masquerades as some well-known protocol, it also has the nasty side-effect of making innovation difficult. This has slowed the deployment of SCTP [RFC4960], since a firewall will often not permit a protocol it does not know even if a user behind it opens the session. When a new protocol or feature is defined, the firewall needs to stop applying that rule, and that can be difficult to make happen.

4.2. Pervasive access control

Another access control model, often called "Role-based", tries to control traffic in flight regardless of the perimeter. Given a rule that equipment located in a given routing domain or with a specific characteristic (such as "student dorms") should not be able to access equipment in another domain or with a specific characteristic (such as "academic records"), it might prevent routing from announcing the second route in the domain of the first, or it might tag individual packets ("I'm from the student dorm") and filter on those tags at

enforcement points throughout network. Such rules can be applied to individuals as well as equipment; in that case, the host needs to tag the traffic, or there must be a reliable correlation between equipment and its user.

One common use of this model is in data centers, in which physical or virtual machines from one tenant (which is not necessarily an "owner" as much as it is a context in which the system is used) might be co-resident with physical or virtual machines from another. Inter-tenant attacks, espionage, and fraud are prevented by enforcing a rule that traffic from systems used by any given tenant is only delivered to other systems used by the same tenant. This might, of course have nuances; under stated circumstances, identified systems or identified users might be able to cross such a boundary.

The major impediment in deployment is complexity. The administration has the option to assign policies for individuals on the basis of their current location (e.g. as the cross-product of people, equipment, and topology), meaning that policies can multiply wildly. The administrator that applies a complex role-based access policy is probably most justly condemned to live in the world he or she has created.

4.3. Intrusion Management: Contract and Reputation filters

The model proposed in Advanced Security for IPv6 CPE [I-D.vyncke-advanced-ipv6-security] could be compared to purchasing an anti-virus software package for one's computer. The proposal is to install a set of filters, perhaps automatically updated, that identify "bad stuff" and make it inaccessible, while not impeding anything else.

It depends on four basic features:

- o A frequently-updated signature-based Intrusion Prevention System which inspects a pre-defined set of protocols at all layers (from layer-3 to layer-7) and uses a vast set of heuristics to detect attacks within one or several flows. Upon detection, the flow is terminated and an event is logged for further optional auditing.
- o A centralized reputation database that scores prefixes for degree of trust. This is unlikely to be on addresses per se, since e.g. temporary addresses [RFC4941] change regularly and frequently.
- o Local correlation of attack-related information, and
- o Global correlation of attacks seen, in a reputation database.

The proposal does not mention anomaly-based intrusion detection, which could be used to detect zero-day attacks and new applications or attacks. This would be an obvious extension.

The comparison to anti-virus software is real; anti-virus software uses similar algorithms, but on API calls or on data exchanged rather than on network traffic, and for identified threats is often effective.

The proposal also has weaknesses:

- o People do not generally maintain anti-virus packages very well, letting contracts expire,
- o Reputation databases have a bad reputation for distributing information which is incorrect, out of date, or compromised by attackers,
- o Anomaly-based analysis identifies changes but is often ineffective in determining whether new application or application behaviors are pernicious (false positives). Someone therefore has to actively decide - a workload the average homeowner might have little patience for, and
- o Signature-based analysis applies to attacks that have been previously identified, and must be updated as new attacks develop. As a result, in a world in which new attacks literally arise daily, the administrative workload can be intense, and reflexive responses like accepting https certificates that are out of date or the download and installation of unsigned software on the assumption that the site administrator is behind are themselves vectors for attack.

Security has to be maintained to be useful, because attacks are maintained.

5. Firewalling Strategies

There is a great deal of tension in firewall policies between two primary goals of networking: the security goal of "block traffic unless it is explicitly allowed" and the networking goal of "trust hosts with new protocols". The two inherently cannot coexist easily in a set of policies for a firewall.

The following subsections discuss the "default deny" and "default allow" security paradigms.

5.1. Blocking Traffic Unless It Is Explicitly Allowed (default deny)

Many networks enforce the so-called "default deny" policy, in which traffic is blocked unless it is explicitly allowed. The rationale for such policy is that it is easier to open "holes" in a firewall to allow specific protocols, than trying to block all protocols that might be employed as an attack vectors; and that a network should only support the protocols it has been explicitly meant to support.

The drawback of this approach is that the security goal of "block traffic unless it is explicitly allowed" prevents useful new applications. This problem has been seen repeatedly over the past decade: a new and useful application protocol is specified, but it cannot get wide adoption because it is blocked by firewalls. The result has been a tendency to try to run new protocols over established applications, particularly over HTTP [RFC3205]. The result is protocols that do not work as well they might if they were designed from scratch.

Worse, the same goal prevents the deployment of useful transports other than TCP, UDP, and ICMP. A conservative firewall that only knows those three transports will block new transports such as SCTP [RFC4960]; this in turn causes the Internet to not be able to grow in a healthy fashion. Many firewalls will also block TCP and UDP options they don't understand, and this has the same unfortunate result.

5.2. Allow Traffic Unless It Is Explicitly Blocked (default allow)

Some networks enforce the so-called "default allow" policy, in which traffic is allowed unless it is explicitly blocked. This policy is usually enforced at perimeters where a comprehensive security policy is not really desirable or possible, but some level of packet filtering is considered appropriate. One common example of such policy could be an ISP blocking TCP port 25 (SMTP), but allowing all other traffic.

When a strict security policy is to be enforced (e.g., at an organizational network's edge), the "default allow" policy tends to be rather inappropriate, since it is usually easier and more effective to identify the traffic that must be allowed through the firewall (and open the necessary "holes" in the firewall) than to identify and block all traffic that may be considered undesirable/inappropriate.

6. Assumptions on IP addresses and Transport Protocol Port Numbers

In a number of scenarios, simple firewall rules have traditionally been specified in terms of the associated IP addresses and transport protocol port numbers. In general, this assumes that the associated IP addresses are stable, and that there is a "well known" transport protocol port number associated with each application.

In the IPv4 world, IP addresses may be considered rather stable. However, IPv6 introduces the concept of "temporary addresses" [RFC4941] which, by definition, change over time. This may prevent the enforcement of filtering policies based on specific IPv6 addresses, or may lead to filtering based on a more coarse granularity (e.g. specific address prefixes, as opposed to specific IPv6 addresses). In some scenarios, from the point of view of enforcing filtering policies, it might be desirable to disable temporary addresses altogether.

For example, an administrator might prefer that a caching DNS server, a secondary DNS server doing zonetransfers, or an SMTP MTA, always employ the same source IPv6 address, as opposed to the temporary addresses that change over time.

The server-side transport protocol port is generally the so-called "well-known port" corresponding to the associated application. While widespread, this practice should probably be considered a kludge/short-cut rather than a "design principle" that can be relied upon for the general case. For example, use of DNS SRV records [RFC2782], or applications such as "portmapper" [Portmap] [RFC1833] might mean that the associated transport protocol port number cannot be assumed to be well-known, but rather needs to be dynamically learned. In other cases, applications may employ (by design) ephemeral port numbers, and there may be no obvious way to dynamically learn the port number being employed. FTP [RFC0959] and SIP [RFC3261] are examples of such applications.

Finally, as a result of widespread packet filtering, many protocols tend to be tunneled employing specific transport-protocol port numbers that are known to be more generally allowed by firewalls, such as TCP port 80 (HTTP). This essentially breaks the assumption that port numbers actually identify the actual application protocol using them.

Some of the so called "next generation" firewalls make fewer assumptions about port numbers, and tend to analyze the application data stream in order to infer the application protocol type, regardless of the well-known port being used. While this may prevent the circumvention of some security controls, it also implies Deep

Packet Inspection (DPI), and therefore there are a number of associated considerations, both in terms of introduced attack vectors and other possibilities for evasion of security controls (please see Section 9 for further discussion).

7. State Associated with Filtering Rules

There are two main paradigms for packet filtering:

- o Stateless filtering
- o Stateful filtering

Stateless filtering implies that the decision on whether to allow or block a specific traffic entity is based solely on the contents of such entity. One common example of such paradigm is the enforcement of network ingress filtering [RFC2827], in which packets may be blocked based on their IP addresses. Stateless filtering scales well, since there are no state requirements on the filtering device other than that associated with maintaining the filtering rules to be applied to the incoming traffic entities (e.g., packets).

On the other hand, stateful filtering implies that the decision on whether to allow or block a traffic entity is not only based on the contents of such entity, but also on the existence (or lack of) previous state associated with such entity. A common example of such paradigm is a firewall that "allows outbound connection requests and only allows return traffic from the external network" (such as the policy implicitly enforced by most NAT devices). For obvious reasons, the firewall needs to maintain state in order to be able to enforce such policies; that is, the firewall may need to keep track of all on-going communication instances, possibly applying timeouts and garbage collection on the associated state table.

Stateful filtering tends to allow more powerful packet filtering, at the expense of increased state. Thus, stateful filtering may be desirable when trying to perform deep packet inspection, but may be undesirable when the firewall is meant to block some Denial of Service attacks, since the firewall itself may become "the weakest link in the chain". Typically, the higher the firewall operates in the network stack, the more state will be required associated. For example, in order for a firewall to enforce a filtering policy based on application-layer request types, the firewall will need to enforce its filtering policy on the application-layer protocol stream, thus implying the need to perform layer-3 and layer-4 reassembly, etc.

When stateful packet filtering is warranted, its associated security implications should be considered. For example, an administrator may

want to enforce traffic filtering to mitigate denial of service attacks; however, when enforcement of such filtering implies increased state at the firewall, the firewall itself may become the easiest target for performing a denial of service attack.

8. Enforcing Protocol Syntax at the Firewall

Some firewalls try to enforce the protocol syntax by checking that only traffic complying with existing protocol definitions is allowed. While this can have useful side-effects, such as preventing the aforementioned traffic from triggering pathological behavior at the target system, it also has the nasty side-effect of making innovation difficult. For example, one of the issues in the deployment of Explicit Congestion Notification [RFC3168] has been that common firewalls often inspect reserved/unused bits and require them to be set to zero to close covert channels. Another example is the plethora of filtering rules applied to DNS traffic [DNS-FILTERING]. When a new protocol or feature is defined, the firewall needs to stop applying that rule, and that can be difficult to make happen.

NOTE:

A somewhat related concept is that of traffic normalization (or "scrubbing"), in which the filtering device can "normalize" traffic by e.g. clearing bits that are expected to be cleared, changing some protocol fields such that they are within "normal" ranges, etc. (see e.g. the discussion of "traffic normalization" in [OpenBSD-PF]). While this can have the useful effect of blocking DoS attacks to sloppy implementations that do not enforce sanity checks on the received packets, it also has the nasty side-effect of making innovation difficult, or even breaking deployed protocols. For example, some firewalls are known to enforce a default packet normalization policy that clears the TCP URG bit, as a result of the TCP urgent mechanism being associated with some popular DoS attacks. Widespread deployment of such firewalls has essentially rendered the TCP urgent mechanism unusable, leading to its eventual formal deprecation in [RFC6093].

We note that, as per our definition of "firewall" in Section 2, "traffic normalization" is not considered a firewall function.

9. Performing Deep Packet Inspection

While filtering packets based on the layer-3 protocol header fields is rather simple and straight-forward, performing enforcing a filtering policy at upper layer protocols can be a challenging task.

For example, IP fragmentation may make this task quite challenging, since even the very layer-4 protocol header could be present in a

non-first fragment. In a similar vein, IPv6 extension headers may represent a challenge for a filtering device, since they can result in long IPv6 extension header chains [RFC7112] [I-D.gont-v6ops-ipv6-ehs-packet-drops].

This problem is exacerbated as one tries to filter packets based on upper layer protocol contents, since many of such protocols implement some form of fragmentation/segmentation and reassembly. In many cases, the reassembly process could possibly lead to different results, and this may be exploited by attackers for circumventing security controls [Ptacek1998] [RFC6274].

In general, the upper in the protocol stack that a filtering policy is to be enforced, the more complex the task becomes: an attacker has more opportunities for obfuscation, ranging from e.g. ambiguities in IP and/or TCP reassembly, to e.g. application-layer obfuscation (such as HTTP URL obfuscation or JavaScript bytecode obfuscation). This usually implies that, in order to reliably enforce a filtering policy, more state is required on the firewall; and the considerations in Section 7 should be evaluated.

10. IANA Considerations

This memo asks the IANA for no new parameters. It can before publication as an RFC by the RFC Editor.

11. Security Considerations

This documents recognizes the role of firewalls in network security, and discusses a number of considerations associated with firewalls which may be of use when designing or deploying firewalls. This document, by itself, does not introduce any security implications.

12. Acknowledgements

The authors would like to thank (in alphabetical order) Fleming Andraeson, Mark Andrews, Lee Howard, Joel Jaeggli, Al Morton, Eric Vyncke and James Woodyatt, for providing valuable comments on earlier versions of this document.

This document is based on [I-D.ietf-opsawg-firewalls-00] authored by Fred Baker, and [I-D.ietf-opsawg-firewalls-01] authored by Paul Hoffman.

13. References

13.1. Normative References

- [RFC1833] Srinivasan, R., "Binding Protocols for ONC RPC Version 2", RFC 1833, DOI 10.17487/RFC1833, August 1995, <<http://www.rfc-editor.org/info/rfc1833>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3205] Moore, K., "On the use of HTTP as a Substrate", BCP 56, RFC 3205, DOI 10.17487/RFC3205, February 2002, <<http://www.rfc-editor.org/info/rfc3205>>.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, DOI 10.17487/RFC4941, September 2007, <<http://www.rfc-editor.org/info/rfc4941>>.
- [RFC6887] Wing, D., Ed., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, DOI 10.17487/RFC6887, April 2013, <<http://www.rfc-editor.org/info/rfc6887>>.
- [RFC7112] Gont, F., Manral, V., and R. Bonica, "Implications of Oversized IPv6 Header Chains", RFC 7112, DOI 10.17487/RFC7112, January 2014, <<http://www.rfc-editor.org/info/rfc7112>>.

13.2. Informative References

- [DNS-FILTERING]
Andrews, M., "On Firewalls in Internet Security (Fwd: New Version Notification for draft-gont-opsawg-firewalls-analysis-00.txt)", post to the OPSAWG mailing-list, Message-Id: <20151012002551.8F7CD3A2FFD8@rock.dv.isc.org>, 2015, <<https://mailarchive.ietf.org/arch/msg/opsawg/2YQl6xBz6jtMyIkyAx59U-oPmPQ>>.
- [I-D.gont-v6ops-ipv6-ehs-packet-drops]
Gont, F., Hilliard, N., Doering, G., LIU, S., and W. Kumari, "Operational Implications of IPv6 Packets with Extension Headers", draft-gont-v6ops-ipv6-ehs-packet-drops-02 (work in progress), February 2016.

- [I-D.ietf-opsawg-firewalls-00]
Baker, F., "On Firewalls in Internet Security", draft-ietf-opsawg-firewalls-00 (work in progress), June 2012.
- [I-D.ietf-opsawg-firewalls-01]
Baker, F. and P. Hoffman, "On Firewalls in Internet Security", draft-ietf-opsawg-firewalls-01 (work in progress), October 2012.
- [I-D.vyncke-advanced-ipv6-security]
Vyncke, E., Yourtchenko, A., and M. Townsley, "Advanced Security for IPv6 CPE", draft-vyncke-advanced-ipv6-security-03 (work in progress), October 2011.
- [OpenBSD-PF]
OpenBSD, , "pf(4) manual page: pf -- packet filter", 2015, <<http://www.openbsd.org/cgi-bin/man.cgi/OpenBSD-current/man4/pf.4&query=pf>>.
- [Portmap] Wikipedia, , "Portmap", 2014, <<https://en.wikipedia.org/wiki/Portmap>>.
- [Ptacek1998]
Ptacek, T. and T. Newsham, "Insertion, Evasion and Denial of Service: Eluding Network Intrusion Detection", 1998, <<http://www.aciri.org/vern/Ptacek-Newsham-Evasion-98.ps>>.
- [RFC0959] Postel, J. and J. Reynolds, "File Transfer Protocol", STD 9, RFC 959, DOI 10.17487/RFC0959, October 1985, <<http://www.rfc-editor.org/info/rfc959>>.
- [RFC2782] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2782, DOI 10.17487/RFC2782, February 2000, <<http://www.rfc-editor.org/info/rfc2782>>.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <<http://www.rfc-editor.org/info/rfc2827>>.
- [RFC2993] Hain, T., "Architectural Implications of NAT", RFC 2993, DOI 10.17487/RFC2993, November 2000, <<http://www.rfc-editor.org/info/rfc2993>>.

- [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, DOI 10.17487/RFC3168, September 2001, <<http://www.rfc-editor.org/info/rfc3168>>.
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, DOI 10.17487/RFC3261, June 2002, <<http://www.rfc-editor.org/info/rfc3261>>.
- [RFC4960] Stewart, R., Ed., "Stream Control Transmission Protocol", RFC 4960, DOI 10.17487/RFC4960, September 2007, <<http://www.rfc-editor.org/info/rfc4960>>.
- [RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", RFC 5575, DOI 10.17487/RFC5575, August 2009, <<http://www.rfc-editor.org/info/rfc5575>>.
- [RFC6092] Woodyatt, J., Ed., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", RFC 6092, DOI 10.17487/RFC6092, January 2011, <<http://www.rfc-editor.org/info/rfc6092>>.
- [RFC6093] Gont, F. and A. Yourtchenko, "On the Implementation of the TCP Urgent Mechanism", RFC 6093, DOI 10.17487/RFC6093, January 2011, <<http://www.rfc-editor.org/info/rfc6093>>.
- [RFC6274] Gont, F., "Security Assessment of the Internet Protocol Version 4", RFC 6274, DOI 10.17487/RFC6274, July 2011, <<http://www.rfc-editor.org/info/rfc6274>>.
- [Saltzer] Saltzer, J., Reed, D., and D. Clark, "End-to-end arguments in system design", ACM Transactions on Computer Systems (TOCS) v.2 n.4, p277-288, Nov 1984.

Authors' Addresses

Fernando Gont
SI6 Networks / UTN-FRH
Evaristo Carriego 2644
Haedo, Provincia de Buenos Aires 1706
Argentina

Phone: +54 11 4650 8472
Email: fgont@si6networks.com
URI: <http://www.si6networks.com>

Fred Baker
Cisco Systems
Santa Barbara, California 93117
USA

Email: fred@cisco.com