

Network (In) Security through IP Packet Filtering

Dhruvika Sharma, Aastha Sharma

Student, Information Technology, Maharishi Dayanand University

New Delhi, Delhi, India

Abstract- Network security has become more important to personal computer users, organizations, and the military. With the advent of the internet, security became a major concern and the history of security allows a better understanding of the emergence of security technology. The internet structure itself allowed for many security threats to occur. The architecture of the internet, when modified can reduce the possible attacks that can be sent across the network. Knowing the attack methods, allows for the appropriate security to emerge. Many businesses secure themselves from the internet by means of firewalls and encryption mechanisms. The businesses create an “intranet” to remain connected to the internet but secured from possible threats. Used properly, packet filtering is a useful tool for the security-conscious network administrator, but its effective use requires a thorough understanding of its capabilities and weaknesses, and of the quirks of the particular protocols that filters are being applied to. This paper examines the utility of IP packet filtering as a network security measure, briefly contrasts IP packet filtering to alternative network security approaches such as application-level gateways, describes what packet filters might examine in each packet, and describes the characteristics of common application protocols as they relate to packet filtering. There are currently two fundamentally different networks, data networks and synchronous network comprised of switches. The internet is considered a data network. Since the current data network consists of computer-based routers, information can be obtained by special programs, such as “Trojan horses,” planted in the routers. The synchronous network that consists of switches does not buffer data and therefore are not threatened by attackers. That is why security is emphasized in data networks, such as the internet, and other networks that link to the internet.

Index Terms- packet filtering works , packet examples , packet filtering caveats and packet filtering rules.

I. INTRODUCTION

The Internet as well as local networks is expanding at a tremendous speed. This one way helps to improve the quality and convenience of the human life but on the other hand provides a platform for network criminals and hackers. An assumption made throughout is that a site administrator is generally more interested in keeping outsiders out than in trying to police insiders, and that the goal is to keep outsiders from breaking in and insiders from accidentally exposing valuable data or services, not to prevent insiders from intentionally and maliciously subverting security measures. This paper does not consider military-grade “secure IP” implementations (those that implement the “IP security options” that may be specified in IP packet headers) and related issues; it is limited to what is commonly available for sale to the general public.

Packet filtering may be used as a mechanism to implement a wide variety of network security policies. The primary goal of these policies is generally to prevent unauthorized network access without hindering authorized network access; the definitions of “unauthorized access” and “authorized access” vary widely from one organization to another. A secondary goal is often that the mechanisms be transparent in terms of performance, user awareness, and application awareness of the security measures. Another secondary goal is often that the mechanisms used be simple to configure and maintain, thus increasing the likelihood that the policy will be correctly and completely implemented; in the words of Bill Cheswick of AT&T Bell Laboratories, “Complex

security isn't". Packet filtering is a mechanism which can, to a greater or lesser extent, fulfil all these goals, but only through thorough understanding of its strengths and weaknesses and careful application of its capabilities.

To lessen the vulnerability of the computer to the network there are many products available. These tools are encryption, authentication mechanisms, intrusion-detection, security management and firewalls. Businesses throughout the world are using a combination of some of these tools. "Intranets" are both connected to the internet and reasonably protected from it. The internet architecture itself leads to vulnerabilities in the network. Understanding the security issues of the internet greatly assists in developing new security technologies and approaches for networks with internet access and internet security itself.

Several factors complicate implementation of these policies using packet filtering, including asymmetric access requirements, differing requirements for various internal and external groups of machines, and the varying characteristics of the particular protocols, services, and implementations of these protocols and services that the filters are to be applied to. Asymmetric access requirements usually arise when an organization desires that its internal systems have more access to external systems than vice versa. Differing requirements arise when an organization desires that some groups of machines have different network access privileges than other groups of machines (for instance, the organization may feel that a particular subnet is more secure than standard, and thus can safely take advantage of expanded network access, or they may feel that a particular subnet is especially valuable, and thus its exposure to the external network should be as limited as possible). Alternatively, an organization may desire to allow more or less network access to some specific group of external machines than to the rest of the external world (for instance, a company might want to extend greater network access than usual to a key client with whom they are collaborating, and less network access than usual to a local university

which is known to be the source of repeated cracker attacks). If a suitably modified version of an application is not available for a given internal host (a modified TELNET client for a personal computer, for instance), that internal host's users are simply out of luck and are unable to reach the past the application gateway.

II. HOW PACKET FILTERING WORKS

2.1. What packet filters base their decisions on

Current IP packet filtering implementations all operate in the same basic fashion; they parse the headers of a packet and then apply rules from a simple rule base to determine whether to route or drop† the packet. Generally, the header fields that are available to the filter are packet type (TCP, UDP, etc.), source IP address, destination IP address, and destination TCP/UDP port. For some reason, the source TCP/UDP port is often *not* one of the available fields; this is a significant deficiency discussed in detail in Section 4.2.

In addition to the information contained in the headers, many filtering implementations also allow the administrator to specify rules based on which router interface the packet is destined to go out on, and some allow rules based on which interface the packet came in on. Being able to specify filters on both inbound and outbound† interfaces allows you significant control over where the router appears in the filtering scheme (whether it is "inside" or "outside" your packet filtering "fence"), and is very convenient (if not essential) for useful filtering on routers with more than two interfaces. If certain packets can be dropped using inbound filters on a given interface, those packets don't have to be mentioned in the outbound filters on all the other interfaces; this simplifies the filtering specifications. Further, some filters that an administrator would like to be able to implement require knowledge of which interface a packet came in on; for instance, the administrator may wish to drop all packets coming inbound from the external interface that claim to be from an internal host, in order to guard against attacks from the outside world that use faked internal source addresses.

Some routers with very rudimentary packet filtering capabilities don't parse the headers, but instead require the administrator to specify byte ranges within the header to examine, and the patterns to look for in those ranges. This is almost useless, because it requires the administrator to have a very detailed understanding of the structure of an IP packet. It is totally unworkable for packets using IP option fields within the IP header, which cause the location of the beginning of the higher-level TCP or UDP headers to vary; this variation makes it very difficult for the administrator to find and examine the TCP or UDP port information.

2.2. How packet filtering rules are specified

Generally, the filtering rules are expressed as a table of conditions and actions that are applied in a certain order until a decision to route or drop the packet is reached. When a particular packet meets all the conditions specified in a given row of the table, the action specified in that row (whether to route or drop the packet) is carried out; in some filtering implementations [Mogul89], the action can also indicate whether or not to notify the sender that the packet has been dropped (through an ICMP message), and whether or not to log the packet and the action taken on it. Some systems apply the rules in the sequence specified by the administrator until they find a rule that applies [Mogul89][Cisco90], which determines whether to drop or route the packet. Others enforce a particular order of rule application based on the criteria in the rules, such as source and destination address, regardless of the order in which the rules were specified by the administrator. Some, for instance, apply filtering rules in the same order as

outing table entries; that is, they apply rules referring to more specific addresses (such as rules pertaining to specific hosts) before rules with less specific addresses (such as rules pertaining to whole subnets and networks) [CHS91][Telebit92a]. The more complex the way in which the router reorders rules, the more difficult it is for the administrator to understand the rules and their application; routers which apply rules in the order specified by the administrator, without reordering the rules, are easier for an administrator to understand and

configure, and therefore more likely to yield correct and complete filter sets.

2.3. A packet filtering example

For example, consider this scenario. The network administrator of a company with Class B network 123.45 wishes to disallow access from the Internet to his network in general (123.45.0.0/16)[†]. The administrator has a special subnet in his network (123.45.6.0/24) that is used in a collaborative project with a local university which has class B network 135.79; he wishes to permit access to the special subnet (123.45.6.0/24) from all subnets of the university (135.79.0.0/16). Finally, he wishes to deny access (except to the subnet that is open to the whole university) from a specific subnet (135.79.99.0/24) at the university, because the subnet is known to be insecure and a haven for crackers. For simplicity, we will consider only packets flowing from the university to the corporation; symmetric rules (reversing the SrcAddr and DstAddr in each of the rules below) would need to be added to deal with packets from the corporation to the university. Rule C is the "default" rule, which specifies what happens if none of the other rules apply.

Rule	SrcAddr	DstAddr	Action
A	135.79.0.0/16	123.45.6.0/24	permit
B	135.79.99.0/24	123.45.0.0/16	deny
C	0.0.0.0/0	0.0.0.0/0	deny

Consider these "sample" packets, their desired treatment under the policy outlined above, and their treatment depending on whether the rules above are applied in order "ABC" or "BAC".

Packet	SrcAddr	DstAddr	Desired Action
1	135.79.99.1	123.45.1.1	deny permit
2	135.79.99.1	123.45.6.1	
3	135.79.1.1	123.45.6.1	permit deny
4	135.79.1.1	123.45.1.1	

A router that applies the rules in the order ABC will achieve the desired results: packets from the "hacker haven" subnet at the university to the company network in general (such as packet 1 above) will be denied (by rule B), packets from the university "hacker haven" subnet at the university to the company's collaboration subnet (such as packet 2 above) will be permitted (by rule A), packets from the

university's general network to the company's "open" subnet (such as packet 3 above) will be permitted (by rule A), and packets from the university's general network to the company's general network (such as packet 4 above) will be denied (by rule C).

If, however, the router reorders the rules by sorting them into order by number of significant bits in the source address then number of significant bits in the destination address, the same set of rules will be applied in the order BAC. If the rules are applied in the order BAC, packet 2 will be denied, when we want it to be permitted.

2.4. Packet filtering caveats

2.4.1. Complexity of packet filtering specifications

In fact, there's a subtle error in this example that illustrates how difficult it is to correctly set up filters using such low-level specifications. Rule B above, which appears to restrict access from the "hacker haven" net, is actually superfluous and unnecessary, and is the cause of the incorrect denial of packet 2 if the rules are applied in the order BAC. If you remove rule B, both types of routers (those that apply rules in the order specified, and those that reorder rules by number of significant bits in source or destination addresses) will process the rules in the order AC. When processed in that order, the result table becomes:

Packet	SrcAddr	DstAddr
1	135.79.99.1	123.45.1.1
2	135.79.99.1	123.45.6.1
3	135.79.1.1	123.45.6.1
4	135.79.1.1	123.45.1.1

There are two points here. First, correctly specifying filters is difficult. Second, reordering filtering rules makes correctly specifying filters even more difficult, by turning a filter set that works (even if it's in fact overspecified) if evaluated in the order given into a filter set that doesn't work.

Even though the example presented above is a relatively simple application of packet filtering, most administrators will probably read through it several times before they feel they understand what is going on. Consider that the more difficult the rules are to comprehend, the

less likely the rules will be correct and complete. The way in which filtering rules must be specified and the order in which they are applied are key determinants of how useful and powerful a given router's filtering capabilities are. Most implementations require the administrator to specify filters in ways which make the filters easy for the router to parse and apply, but make them very difficult for the administrator to comprehend and consider.

2.4.2. Reliance on accurate IP source addresses

Most filtering implementations, of necessity, rely on the accuracy of IP source addresses to make filtering decisions. IP source addresses can easily be faked, however, as discussed in [Bellovin89], [Kent89], [Bellovin92a], and [Bellovin92b]. This is a particular case where being able to filter inbound packets is useful. If a packet that appears to be from one internal machine to another internal machine comes in over the link from the outside world, you should be mighty suspicious. If your router can be told to drop such packets using inbound filters on the external interface, your filtering specifications for internal interfaces can be made both much simpler and more secure.

2.4.3. Dangers of IP source routing

Another IP feature ripe for potential abuse is IP source routing. Essentially, an IP packet with source routing information included tells the router how to get to the packet, rather than letting the router decide for themselves. An attacker could use this to their advantage [Bellovin89]. Unless you have a specific need to allow packets with IP source routes between your internal network and the outside world, it's probably a good idea for your router to ignore IP source route instructions; whether source routing can be disabled, whether it is enabled or disabled by default, and how to disable it vary from vendor to vendor.

2.4.4. Complications due to IP fragmentation

Yet another complication to packet filtering is IP packet fragmentation. IP supports the notion that any router along a packet's path may "fragment" that packet into several smaller

packets, to accommodate the limitations of underlying media, to be reassembled into the original IP packet at the destination. For instance, an FDDI frame is much larger than an Ethernet frame; a router between an FDDI ring and an Ethernet may need to split an IP packet that fit in a single FDDI frame into multiple fragments that fit into the smaller Ethernet frames. The problem with this, from a packet filtering point of view, is that only the first of the IP fragments has the higher-level protocol (TCP or UDP) headers from the original packet, which may be necessary to make a filtering decision concerning the fragment. Different filtering implementations take a variety of responses to this situation. Some apply filters only to the first fragment (which contains the necessary higher-level protocol headers), and simply route the rest, on the assumption that if the first fragment is dropped by the filters, the rest of the fragments can't be reassembled into a full packet, and will cause no harm [CHS91]. Others keep a cache of recently-seen first fragments and the filtering decision that was reached, and look up non-first fragments in this cache in order to apply the same decision [Mogul89]. In particular, it is dangerous to suppress only the first fragment of outbound packets; you might be leaking valuable data in the non-first fragments that are routed on out.

III. FILTERING-RELATED CHARACTERISTICS OF APPLICATION PROTOCOLS

Each application protocol has its own particular characteristics that relate to IP packet filtering, that may or may not differ from other protocols. Particular implementations of a given protocol also have their own characteristics that are not a result of the protocol per se, but a result of design decisions made by the implementors. Since these implementation characteristics are not covered in the specification of the protocol (though they aren't counter to the specification), they are likely to vary between different implementations of the same protocol, and might change even within a given implementation as that implementation evolves. These characteristics include what port a server uses, what port a client uses, whether the service is

typically offered over UDP or TCP or both, and so forth. An understanding of these characteristics is essential for setting up effective filters to allow, disallow, or limit the use of these protocols. Appendix A discusses in detail the filtering-related characteristics of several common protocols.

III.1. "Random" ports aren't really random

Although implementations of various protocols might appear to use a "random" ports for the client end and a well-known port for the server end, the ports chosen for the client end used are usually not totally random. While not explicitly supported by the RFCs, systems based on BSD UNIX usually reserve ports below 1024 for use by "privileged" processes, and allow only processes running as root to bind to those ports; conversely, non-privileged processes must use ports at or above 1024. Further, if a program doesn't request a particular port, it is often simply assigned the port after the last one assigned; if the last port assigned was 5150, the next one assigned will probably be 5151.

III.2.Privileged versus non-privileged ports

The distinction between "privileged" and "non-privileged" ports (those below 1024 and at or above 1024, respectively) is found throughout BSD-based systems (and other systems that draw from a BSD background; keep in mind that almost all UNIX IP networking, including SysV IP networking, draws heavily from the original BSD network implementation). This distinction is not codified in the RFCs, and is therefore best regarded as a widely used convention, but not as a standard. Nonetheless, if you're protecting UNIX systems, the convention can be a useful one. You can, for instance, generally forbid all inbound connections to ports below 1024, and then open up specific exceptions for specific services that you wish to enable the outside world to use, such as SMTP, TELNET, or FTP; to allow the "return" packets for connections to such services, you allow all packets to external destination ports at or above 1024.

While it would simplify filtering if all services were offered on ports below 1024 and all clients used ports at or above 1024, many

vulnerable services (such as X, OpenWindows, and a number of database servers) use server ports at or above 1024, and several vulnerable clients (such as the Berkeley r* programs) use client ports below 1024. These should be carefully excepted from the "allow all packets to destination ports at or above 1024" type of rules that allow return packets for outbound services.

IV. PROBLEMS WITH CURRENT PACKET FILTERING IMPLEMENTATIONS

IP packet filtering, while a useful network security tool, is not a panacea, particularly in the form in which it is currently implemented by many vendors. Problems with many current implementations include complexity of configuration and administration, omission of the source UDP/TCP port from the fields that filtering can be based on, unexpected interactions between "unrelated" parts of the filter rule set, cumbersome filter specifications forced by simple specification mechanisms, a lack of testing and debugging tools, and an inability to deal effectively with RPC-based protocols such as YP/NIS and NFS.

IV.1. Filters are difficult to configure

The first problem with many current IP packet filtering implementations as network security mechanisms is that the filtering is usually very difficult to configure, modify, maintain, and test, leaving the administrator with little confidence that the filters are correctly and completely specified. The simple syntax used in many filtering implementations makes life easy for the router (it's easy for the router to parse the filter specifications, and fast for the router to apply them), but difficult for the administrator (it's like programming in assembly language). Instead of being able to use high-level language abstractions ("if this and that and not something-else then permit else deny"), the administrator is forced to produce a tabular representation of rules; the desired behavior may or may not map well on to such a representation.

Administrators often consider networking activity in terms of "connections", while packet filtering, by definition, is concerned with the packets making up a connection. An administrator might think in terms of "an

inbound SMTP connection", but this must be translated into at least two filtering rules (one for the inbound packets from the client to the server, and one for the outbound packets from the server back to the client) in a table-driven filtering implementation. The concept of a connection is applied even when considering a connectionless protocol such as UDP or ICMP; for instance, administrators speak of "NFS connections" and "DNS connections". This mismatch between the abstractions used by many administrators and the mechanisms provided by many filtering implementations contributes to the difficulty of correctly and completely specifying packet filters.

4.2. Special handling of start-of-connection packets is impossible

Note that the even the above filters with source port still don't protect your servers living at or above port 1024 from an attack launched from port 25 on an external machine (which is certainly possible if the person making the attack controls the machine the attack is coming from); rules C and D will allow this. One way to defeat this type of attack is to suppress TCP start-of-connection packets (packets with the TCP "SYN" flag set) in rule C; at least one filter implementation provides a mechanism for stating that rules apply *only* to packets in "established" connections (those packets without the SYN bit set) [Cisco90].

Unfortunately, UDP sessions are "connectionless", so there is never a "start-of-connection" packet that can be suppressed in a UDP session. A solution for UDP is often to disallow UDP entirely except for a specific exception for DNS. This exception for DNS can generally be made safely even with a filtering implementation that ignores source port, because of a quirk in the most common DNS implementation. The quirk causes DNS server-to-server queries made over UDP to always use port 53 at both ends of the connection, rather than a random port at one end. Disallowing UDP except for DNS also allows you to avoid most of the problems with filtering RPC-based services (since most RPC services are UDP based) that are discussed in Section 4.6.

4.3. Tabular filtering rule structures are too cumbersome

While tabular rule structures such as those shown above are relatively easy and thus efficient for the router to parse and apply, they rapidly become too cumbersome for the administrator to use to specify complex independent filtering requirements. Even simple applications of these cumbersome syntaxes are difficult, and often have unintended and undesired side effects, as demonstrated in Section 4.2.

4.4. Testing and monitoring filters is difficult

With many router products, the beleaguered administrator's life is further complicated by a lack of built-in mechanisms to test the filter set or to monitor its performance in action. This makes it very difficult to debug and validate filtering rule sets, or to modify existing rule sets; the administrator always has to wonder if the filtering rules are really accomplishing what was intended, or if the rule set has some inadvertent hole in it that the administrator has somehow overlooked.

4.5. RPC is very difficult to filter effectively

Finally, RPC-based protocols offer a special challenge, since they don't reliably appear on a given UDP or TCP port number. The only RPC-related service that is guaranteed to be at a certain port is the "portmapper" service. Portmapper maps an RPC service number (which is a 32-bit number assigned by Sun Microsystems to each individual RPC service, including services created by users and other vendors) to the particular TCP or UDP port number (which are much smaller 16-bit numbers) that the service is currently using on the particular machine being queried. When an RPC-based service starts up, it registers with the portmapper to announce what port it is living at; the portmapper then passes this info along to anyone who requests it.

The portmapper isn't required in order to establish an RPC connection, except to determine exactly which port to establish the connection to; if you know (or can guess) which port to establish the connection to, you can bypass the portmapper altogether. What port a given RPC protocol (such as YP/NIS, NFS, or any of a

number of others) ends up using is random enough that the administrator can't effectively specify filters for it (at least, not without risking the inadvertent filtering of something else that happened to end up on the same port the administrator thought an RPC-based service *might* end up at), but not so random that an attacker can't easily "guess" where a given protocol lives. Even if they can't or don't guess, a systematic search of the entire port number space for the RPC service they're interested in attacking is not that difficult. Since RPC-based services might be on any port, the filtering implementation has no sure way of recognizing what is and what isn't RPC; as far as the router is concerned, it's all just UDP or TCP traffic.

Two fortuitous characteristics of most RPC-based services can be used to save us from this morass, however. First, most RPC-based services are offered as only on UDP ports; we can simply drop UDP packets altogether except for DNS, as described above. Second, almost all of those that are offered on TCP ports use ports below 1024, which can be protected by an "deny all ports below 1024 except specific services like SMTP" type of filter, such as shown in the example in Section 4.2.

V. POSSIBLE SOLUTIONS FOR CURRENT PACKET FILTERING PROBLEMS

V.1. Improve filter specification syntax

The major improvement that could be made to many vendor packet filtering implementations would be to provide better filter specification mechanisms. The administrator should be able to specify rules in a form that makes sense to the administrator (such as a propositional logic syntax), not necessarily a form that is efficient for the router to process; the router can then convert the rules from the high-level form to a form amenable to efficient processing. One possibility might be the creation of a "filter compiler" that accepts filters in a high-level syntax that was convenient for the administrator, and emits a "compiled" filter list that is acceptable to the router.

Addressing the conceptual mismatch between administrators, who think in terms of

connections, and routers, which operate in terms of the packets making up those connections, as discussed in Section 4.1, might also prove valuable.

V.2. Make all relevant header fields available as filtering criteria

The administrator should be able to specify all relevant header fields, particularly including TCP/UDP source port (which is currently often omitted from many filtering implementations), as filter criteria. Until this key feature is provided, it will be difficult or impossible to effectively use filtering in certain common situations, as demonstrated in the example in Section 4.2. The administrator should also be able to specify whether a filter rule should apply only to established TCP connections.

V.3. Allow inbound filters as well as outbound filters

The administrator should be able to specify both inbound and outbound filters on each interface, rather than only outbound filters. This would allow the administrator to position the router either "inside" or "outside" the filtering "fence", as appropriate. It would also allow simpler specification of filters on routers with more than two interfaces by allowing some cases (such as a packet appearing from the outside world that purports to be both to and from internal hosts) to be handled by the inbound set of filters on the external interface, rather than having to duplicate these special cases into the outbound filter set on each internal interface. The desired functionality may not even be possible with only outbound filters; the case of a fake internal-to-internal packet showing up on the external interface, as discussed in Section 2.4.2, can't be detected in an outbound filter set.

V.4. Provide tools for developing, testing, and monitoring filters

Better tools for developing, testing and validating rule sets, perhaps including test suites and automatic test probe generators, would make a big difference in the usability of packet filtering mechanisms. Such an automated test system might well be a part of the "filter compiler" described in Section 5.1.

V.5. Simplify specification of common filters

It would be useful if administrators could specify common filtering cases (for instance, "allow inbound SMTP to this single host") simply, without having to understand the details of the protocols or filtering mechanisms involved.

VI. CONCLUSIONS

Packet filtering is currently a viable and valuable network security tool, but some simple vendor improvements could have a big impact. There are several critical deficiencies that seem to be common to many vendors, such as the inability to consider source TCP/UDP port in filters, that need to be addressed. Other improvements to filter specification mechanisms could greatly simplify the lives of network administrators trying to use packet filtering capabilities, and increase their confidence that their filters are doing what they think they are.

VII. ACKNOWLEDGEMENTS

Thanks to Steve Bellovin and Bill Cheswick of AT&T Bell Laboratories for several lively and fruitful discussions of packet filtering as a network security tool; in particular, I'd like to thank Steve for providing me with prepublication copies of two of his IP security-related papers and of his 1989 article on TCP/IP security problems. Thanks to Ed DeHart of the Computer Emergency Response Team for strongly and repeatedly encouraging me to write this paper after listening to me moan about the issues discussed herein. Thanks to Elizabeth Zwicky of SRI International, Brian Lloyd of Lloyd & Associates, and Steve Bellovin of AT&T Bell Laboratories for reviewing drafts of this paper and providing valuable feedback and suggestions.

REFERENCES

- [Bellovin89]
S. M. Bellovin, "Security Problems in the TCP/IP Protocol Suite"; *Computer Communications Review*, Volume 9, Number 2; April 1989; pp. 32-48.
- [Bellovin92a]
Steven M. Bellovin, "Packets Found on an Internet"; in preparation; 1992.

- [Bellovin92b]
Steven M. Bellovin, "There Be Dragons"; *Proceedings of the Third USENIX UNIX Security Symposium*; Baltimore, MD; September, 1992.
- [Ches90]
Bill Cheswick, "The Design of a Secure Internet Gateway"; *Proceedings of the USENIX Summer 1990 Conference*; Anaheim, CA; June 11-15, 1990; pp. 233-237.
- [CHS91]
Bruce Corbridge, Robert Henig, Charles Slater, "Packet Filtering in an IP Router"; *Proceedings of the Fifth USENIX Large Installation and System Administration Conference (LISA V)*; San Diego, CA; October, 1992; pp. 227-232.
- [Cisco90]
Cisco Systems (Menlo Park, CA); "Gateway System Manual; Software Release 8.2"; 1990.
- [CMQ92]
Smoot Carl-Mitchell and John S. Quarterman, "Building Internet Firewalls"; *UnixWorld*; February, 1992; pp 93-102.
- [Comer91]
Douglas E. Comer, *Internetworking with TCP/IP, Volume I*; Second Edition, 1991; Prentice-Hall, Inc.
- [Kent89]
Stephen Kent, "Comments on 'Security Problems in the TCP/IP Protocol Suite'"; *Computer Communications Review*; July 1989.
- [Mogul89]
Jeffrey C. Mogul, "Simple and Flexible Datagram Access Controls for UNIX-based Gateways"; *Proceedings of the USENIX Summer 1989 Conference*; pp. 203-221.
- [Ranum92]
Marcus J. Ranum, "A Network Firewall"; *Proceedings of the World Conference on System Administration and Security*; July 1992; Washington, D.C.; pp. 153-163.
- [RFC1058]
C. Hedrick, "Routing Information Protocol", Request For Comments 1058; available from the DDN Network Information Center (NIC.DDN.MIL).
- [RFC1340]
J. Reynolds and J. Postel, "Assigned Numbers", Request For Comments 1340; available from the DDN Network Information Center (NIC.DDN.MIL).
- [Telebit92a]
Telebit Corporation (Sunnyvale, CA), "NetBlazer Command Reference"; 1992.
- [Telebit92b]
Telebit Corporation (Sunnyvale, CA), "NetBlazer Version 1.4 Release Notes"; 1992.

Appendix A — Filtering Characteristics of Common IP Protocols

A.1. SMTP

SMTP is provided as a TCP service with the server end of the connection at port 25 and the client end at a random port.

A.2. TELNET

TELNET is provided as a TCP service with the server end of the connection at port 23, and the client end at a random port.

A.3. FTP

FTP is slightly tricky, in that an FTP conversation actually involves two TCP connections in typical UNIX implementations: one for connection for commands, and one for data. The command connection is at port 21 on the server, and the data connection is at port 20 on the server; both connections use random ports on the client side.

A.4. NNTP

NNTP is provided as a TCP service with the server end at port 119, and the client end at a random port.