

Network Working Group
Internet-Draft
Expires: January 8, 2005

J. Touch
USC/ISI
July 10, 2004

Defending TCP Against Spoofing Attacks
draft-touch-tcp-antispoof-00

Status of this Memo

By submitting this Internet-Draft, I certify that any applicable patent or other IPR claims of which I am aware have been disclosed, and any of which I become aware will be disclosed, in accordance with [RFC 3668](#). This document may not be modified, and derivative works of it may not be created.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on January 8, 2005.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

Recent attacks on core Internet infrastructure indicate an increased vulnerability of TCP connections to spurious resets (RSTs). TCP has always been susceptible to such RST spoof attacks, which were indirectly protected by checking that the RST sequence number was inside the current receive window, as well as via the obfuscation of TCP endpoint and port numbers. For pairs of well-known endpoints often over predictable port pairs, such as BGP, increases in the path bandwidth-delay product of a connection have sufficiently increased the receive window space that off-path third parties can guess a

Touch

Expires January 8, 2005

[Page 1]

viable RST sequence number. This document addresses this vulnerability, discussing proposed solutions at the transport level and their inherent challenges, as well as existing network level solutions and the feasibility of their deployment.

Table of Contents

1.	Introduction	3
2.	Background	4
2.1	Recent BGP Attacks Using TCP RSTs	4
2.2	TCP RST Vulnerability	5
2.3	What Changed -- the Ever Opening Window	5
3.	Proposed solutions	9
3.1	Transport Layer	9
3.1.1	TCP MD5 Authentication	9
3.1.2	TCP RST Window Attenuation	9
3.1.3	TCP Timestamp Authentication	10
3.1.4	Other TCP Cookies	11
3.1.5	Other TCP Considerations	11
3.1.6	Other Protocols	12
3.2	Network Layer (IP)	12
4.	Issues	13
4.1	Transport Layer (e.g., TCP)	13
4.2	Network Layer (IP)	14
4.3	Application Layer	15
4.4	Shim Transport/Application Layer	15
4.5	Link Layer	15
4.6	Need for Alternate Security Levels	15
5.	The Need for High-Performance Anonymous Security	17
5.1	Anonymous Keying	17
5.2	Alternatives for Performance	17
6.	Security Considerations	20
7.	Conclusions	21
8.	Acknowledgments	22
9.	References	23
9.1	Normative References	23
9.2	Informative References	24
	Author's Address	25
	Intellectual Property and Copyright Statements	26

Touch

Expires January 8, 2005

[Page 2]

1. Introduction

The Internet infrastructure has recently seen a flurry of attacks on BGP connections between core routers using an attack known for nearly six years [[1](#)][[2](#)]. These connections, typically using TCP, can be susceptible to off-path (non man-in-the-middle) third-party reset (RST) segments, which terminate the TCP connection. BGP routers react to a terminated TCP connection in various ways, ranging from restarting the connection to deciding that the other router is unreachable and thus flushing the BGP routes. This sort of attack affects other protocols besides BGP, involving any long-lived connection between well-known endpoints. The impact on Internet infrastructure has been substantial (esp. for the BGP case), and warrants immediate attention.

TCP, like many other protocols, has been susceptible to off-path third-party attacks. Such attacks rely on the increase of commodity platforms supporting public access to previously privileged resources, such as root-level access. Given such access, it is trivial for anyone to generate a packet with any header desired. This, coupled with the lack of sufficient ingress filtering to drop such spoofed traffic, has resulted in an increase in off-path third-party spoofing attacks. As a result, a number of proposed solutions have been developed in collaboration among the Internet research and commercial router communities. The foremost of these modifies TCP processing to defeat off-path third-party spoofs by further limiting viable sequence numbers in the RST segment.

Such modifications are, at best, temporary patches to the ubiquitous vulnerability to spoofing attacks. The obvious solution to spoofing is to validate the segments of a connection, either at the transport level or the network level. TCP with MD5 extensions already provides this authentication at the transport level, and IPsec already intends to provide authentication at a network level, but in both cases their deployment overhead can be prohibitive. E.g., it is not feasible for BGP routers to be configured with the appropriate certificate authorities of large numbers of peers (for IPsec using IKE), or shared secrets (for IPsec in shared-secret mode, or TCP/MD5), because many clients may need to be configured rapidly without external assistance. The same is true for other uses of long-lived TCP connections between well-known pairs.

The remainder of this document outlines the attack in detail and describes and compares a variety of solutions, including existing solutions based on TCP/MD5 and IPsec, as well as recently proposed solutions, including modifications to TCP's RST processing [[22](#)], modifications to TCP's timestamp processing [[3](#)], and modifications to IPsec and TCP/MD5 keying [[4](#)].

Touch

Expires January 8, 2005

[Page 3]

2. Background

The recent attacks on BGP have raised the issue of TCP's vulnerability to off-path third-party spoofing attacks [[1](#)]. A number of such attacks have been known for several years, including sending RSTs, SYNs, and even ACKs in an attempt to affect an existing connection or to load down servers. Overall, such attacks are countered by the use of some form of authentication at the network (IPsec), transport (SYN cookies), or other layers. TCP already includes a weak form of such authentication in its check of segment sequence numbers. Increases in the bandwidth-delay product for certain long connections has made sufficiently weakened this authentication in recent weeks, rendering it moot.

2.1 Recent BGP Attacks Using TCP RSTs

BGP represents a particular vulnerability to spoofing attacks. Most TCP connections are protected by multiple levels of obfuscation except at the endpoints of the connection:

- o Both endpoint addresses are usually not well-known; although server addresses are advertised, clients are somewhat anonymous.
- o Both port numbers are usually not well-known; the server's usually is advertised (representing the service), but the client's is typically sufficiently unpredictable to an off-path third-party.
- o Valid sequence number space is not well-known.
- o Connections are relatively short-lived and valid sequence space changes, so any guess of the above information is unlikely to be useful.

BGP represents an exception to the above criteria (though not the only case). Both endpoints are well-known, notably as part of an AS path. The destination port is typically fixed to indicate the BGP service. The source port used by a BGP router is sometimes fixed and advertised to enable firewall configuration; even when not fixed, there are only 65,000 valid source ports which may be exhaustively attacked. Connections are long-lived, and as noted before some BGP implementations interpret successive TCP connection failures as routing failures, discarding the corresponding routing information. As importantly and as will be shown below, the valid sequence number space once thought to provide some protection has been rendered useless by increasing congestion window sizes.

Touch

Expires January 8, 2005

[Page 4]

2.2 TCP RST Vulnerability

TCP has a known vulnerability to third-party spoofed segments. SYN flooding consumes server resources in half-open connections, affecting the server's ability to open new connections. ACK spoofing can cause connections to transmit too much data too quickly, creating network congestion and segment loss, causing connections to slow to a crawl. In the most recent attacks on BGP, RSTs cause connections to be dropped. As noted earlier, some BGP implementations interpret TCP connection termination, or a series of such failures, as a network failure. This causes routers to drop the BGP routing information already exchanged, in addition to inhibiting their ongoing exchanges. The result can affect routing paths throughout the Internet.

The dangerous effects of RSTs on TCP have been known for many years, even when used by the legitimate endpoints of a connection. TCP RSTs cause the receiver to drop all connection state; because the source is not required to maintain a TIME_WAIT state, such a RST can cause premature reuse of address/port pairs, potentially allowing segments from a previous connection to contaminate the data of a new connection, known as TIME_WAIT assassination [5]. In this case, assassination occurs inadvertently as the result of duplicate segments from a legitimate source, and can be avoided by blocking RST processing while in TIME_WAIT. However, assassination can be useful to deliberately reduce the state held at servers; this requires that the source of the RSTs go into TIME_WAIT state to avoid such hazards, and that RSTs are not blocked in the TIME_WAIT state [6].

Firewalls and load balancers, so-called 'middleboxes', sometimes emit RSTs on behalf of transited connections to optimize server performance [7]. This is a 'man in the middle' RST attack, where the RSTs are sent for benign or beneficial intent. There are numerous hazards with such use of RSTs, outlined in that RFC.

2.3 What Changed -- the Ever Opening Window

RSTs represent a hazard to TCP, especially when completely unchecked. Fortunately, there are a number of obfuscation mechanisms that make it difficult for off-path third parties to forge (spoof) valid RSTs, as noted earlier. We have already shown it is easy to learn both endpoint addresses and ports for some protocols, notably BGP. The final obfuscation is the sequence number.

TCP segments include a sequence number which enables out-of-order receiver processing, as well as duplicate detection. The sequence number space is also used to manage congestion, and indicates the index of the next byte to be transmitted or received. For RSTs, this is relevant because legitimate RSTs use the next sequence number in

Touch

Expires January 8, 2005

[Page 5]

the transmitter window, and the receiver checks that incoming RSTs have a sequence number in the expected receive window. Such processing is intended to eliminate duplicate segments (somewhat moot for RSTs, though), and to drop RSTs which were part of previous connections.

TCP uses two window mechanisms, a primary mechanism which uses a space of 32 bits, and a secondary mechanism which scales this window [8][9]. The valid receive window is a fraction, not to exceed approximately half, of this space, or ~2,000,000,000. Under typical use, the majority of TCP connections open to a very small fraction of this space, e.g., 10,000-60,000 (approximately 5-100 segments). On a low-loss path, the window should open to around the path bandwidth-delay product, including buffering delays (assume 1 packet/hop). Many paths in the Internet have end-to-end bandwidths of under 1 Mbps, latencies under 100ms, and are under 15 hops, resulting in fairly small windows as above (under 35,000 bytes). Under these conditions, and further assuming that the initial sequence number is suitably (pseudo-randomly) chosen, a valid guessed sequence number would have odds of 1 in 57,000. Put differently, a blind (non man-in-the-middle) attacker would need to send 57,000 RSTs with suitably spaced sequence number guesses to successfully reset a connection. At 1 Mbps, 57,000 (40 byte) RSTs would take over 50 minutes to transmit, and, as noted earlier, most current connections are fairly brief by comparison.

Recent use of high bandwidth paths of 10 Gbps and result in bandwidth-delay products over 125 MB - approximately 1/10 of TCP's overall window size excluding scale, assuming the receiver allocates sufficient buffering (to be discussed later). Even under networks that are ten times slower (1 Gbps), the active receiver window covers 1/100th of the overall window size. At these speeds, it takes only 10-100 packets, or under 32 microseconds, to correctly guess a valid sequence number and kill a connection. A table of corresponding exposure to various amounts of RSTs is shown below, for various line rates, assuming the more conventional 100ms latencies (though even 100ms is large for BGP cases):

Touch

Expires January 8, 2005

[Page 6]

BW	BW*delay		RSTs needed	Time needed

10 Gbps	125	MB	35	1 us (microsecond)
1 Gbps	12.5	MB	344	110 us
100 Mbps	1.25	MB	3,436	10 ms (millisecond)
10 Mbps	0.125	MB	34,360	1 second
1 Mbps	0.0125	MB	343,598	2 minutes
100 Kbps	0.00125	MB	3,435,974	3 hours

Figure 1: Time needed to kill a connection

This table demonstrates that the effect of bandwidth on the vulnerability is squared; for every increase in bandwidth, there is a linear decrease in the number of sequence number guesses needed, as well as a linear decrease in the time needed to send a set of guesses. Notably, as inter-router link bandwidths approach 1 Mbps, an 'exhaustive' attack becomes practical. Checking that the RST sequence number is somewhere in the valid window ($bw \cdot delay$) out of the overall window (2^{32}) is an insufficient obfuscation.

Note that this table makes a number of assumptions:

1. the overall bandwidth-delay product is relatively fixed
2. traffic losses are negligible (insufficient to affect the congestion window over most of the connection)
3. the receive socket buffers do not limiting the receive window
4. the attack bandwidth is similar to the end-to-end path bandwidth

Of these assumptions, the last two are more notable. The issue of receive socket buffers will be addressed later. The issue of the attack bandwidth is considered reasonable as follows:

1. RSTs are substantially easier to send than data; they can be precomputed and they are smaller than data packets (40 bytes).
2. although susceptible connections use somewhat less ubiquitous high-bandwidth paths, the attack may be distributed, at which point only the ingress link of the attack is the primary limitation
3. for the purposes of the above table, we assume that the ingress at the attack has the same bandwidth as the path, as an approximation

The previous sections discussed the nature of the recent attacks on

Touch

Expires January 8, 2005

[Page 7]

BGP due to the vulnerability of TCP to RST spoofing attacks, due largely to recent increases in the fraction of the TCP window space in use for a single, long-lived connection.

3. Proposed solutions

TCP currently authenticates received RSTs using the address and port pair numbers, and checks that the sequence number is inside the valid receiver window. The previous section demonstrated how TCP has become more vulnerable to RST spoofing attacks due to the increases in the receive window size. There are a number of current and proposed solutions to this vulnerability, all centering on increasing the authentication of received RSTs.

3.1 Transport Layer

The transport layer represents the last place that segments can be authenticated before they affect connection management. TCP has a variety of current and proposed mechanisms to increase the authentication of segments, protecting against both off-path third-party spoofs and man-in-the-middle attacks. SCTP also has mechanisms to authenticate segments.

3.1.1 TCP MD5 Authentication

An extension to TCP supporting MD5 authentication was developed around six years ago specifically to authenticate BGP connections [2]. The extension relies on a pre-shared secret key to authenticate the entire TCP segment, including the data, TCP header, and TCP pseudo-header (certain fields of the IP header). All segments are protected, including RSTs, which are accepted only when their signature matches. This option, although widely deployed in Internet routers, is considered undeployable for widespread use because the need for pre-shared keys. It further is considered computationally expensive for either hosts or routers due to the overhead of MD5 [10][11].

3.1.2 TCP RST Window Attenuation

A recent proposal extends TCP to further constrain received RST to match the expected next sequence number [22]. This restores TCP's resistance to spurious RSTs, effectively limiting the receive window for RSTs to a single number. As a result, an attacker would need to send 2^{32} different packets to correctly guess the sequence number. The extension further modifies the RST receiver to react to incorrectly-numbered RSTs, by sending a zero-length ACK. If the RST source is legitimate, upon receipt of an ACK the closed source would presumably emit a RST with the sequence number matching the ACK, correctly resetting the intended recipient. There are a number of concerns with this proposal, including the platitude "think twice before modifying TCP, then don't" [12], notably because this modification adds arcs to the TCP state diagram (in contrast to

Touch

Expires January 8, 2005

[Page 9]

adding MD5 signatures, which is orthogonal to the state machine altogether). For example, there may be complications between RSTs of different connections between the same pair of endpoints because RSTs flush the TIME-WAIT (as mentioned earlier). Further, this modifies TCP so that under some circumstances a RST causes a reply, in violation of generally accepted practice, if not gentle recommendation. The advantage to this proposal is that it can be deployed incrementally and has benefit to the endpoint on which it is deployed.

A variant of this proposal uses a different value to attenuate the window of viable RSTs. It requires RSTs to carry the initial sequence number rather than the next expected sequence number, i.e., the value negotiated on connection establishment [13]. This proposal has the advantage of using an explicitly negotiated value, but at the cost of changing the behavior of an unmodified endpoint to a currently valid RST. It would thus be more difficult, without additional mechanism, to deploy incrementally.

The most obvious other variant of this proposal involves increasing TCP's window space, rather than decreasing the valid range for RSTs, i.e., increasing the sequence space from 32 bits to 64 bits. This has the equivalent effect - the ratio of the valid sequence numbers for any segment to the overall sequence number space is significantly reduced. The use of the larger space, as with current schemes to establish weak authentication using initial sequence numbers (ISNs), is contingent on using suitably random values for the ISN. Such randomness adds additional complexity to TCP both in specification and implementation, and provides only very weak authentication at best. While there are many reasons to increase the TCP sequence number space, we believe authentication is not one of them. Finally, such a modification is not obviously backward compatible, and would be thus difficult to deploy.

3.1.3 TCP Timestamp Authentication

Another way to authenticate TCP segments is to utilize its timestamp option, using the value as a sort of authentication [3]. This requires that the receiver TCP discard values whose timestamp is outside the accepted window, which is derived from the timestamps of other packets from the same connection. This technique uses an existing TCP option, but also requires modified RST processing and may be difficult to deploy incrementally without further modifications. Additionally, the timestamp value may be easier to guess because it is derived from a predictable value.

Touch

Expires January 8, 2005

[Page 10]

3.1.4 Other TCP Cookies

All of the above techniques are variants of cookies, otherwise nonsensical data whose value is used to validate the packet. In the case of MD5 checksums, the cookie is computed based on a shared secret. Even a signature can be guessed, and presents a 1 in $2^{(\text{cookie length})}$ probability of attack anyway. The primary difference is that MD5 signatures are effectively one-time cookies, not predictable based on man-in-the-middle snooping. Window attenuation sequence numbers can be guessed by snooping the sequence number of current packets, and timestamps can may be guessed even more remotely. These variants of cookies are similar in spirit to TCP SYN cookies, again patching a vulnerability to off-path third-party spoofing attacks based on a (fairly weak, excepting MD5) form of authentication. Another form of cookie is the source port itself, which can be randomized but provides only 16 bits of protection (65,000 combinations), which may be exhaustively attacked. This can be combined with destination port randomization as well, but that would require a separate coordination mechanism (so both parties know which ports to use), which is equivalent to (and as infeasible for large-scale deployments as) exchanging a shared secret.

3.1.5 Other TCP Considerations

The analysis of the potential for RST spoofing above assumes that the receive window opens to the maximum extent suggested by the bandwidth-delay product of the end-to-end path, and that the window opens to an appreciable fraction of the overall sequence number space. As noted earlier, for most common cases, connections are too brief or over bandwidths too low to for such a large window to occur. Expanding TCP's sequence number space is a direct way to further avoid such vulnerability, even for long connections over emerging bandwidths.

Finally, it is often sufficient for the endpoint to limit the receive window in other ways, notably using 'socket options'. If the receive socket buffer is limited, e.g., to the ubiquitous default of 65KB, the receive window cannot grow to vulnerable sizes even for very long connections over very high bandwidths. The consequence is lower sustained throughput, where only one window's worth of data per round trip time (RTT) is exchanged. Although this will keep the connection open longer, it also reduces the receive window; for long-lived connections with continuous sourced data, this may continue to present an attack opportunity, albeit a sparse and slow-moving target. For the most recent case where BGP data is being exchanged between Internet routers, the data is bursty and the aggregate traffic is small (i.e., unlikely to cover a substantial portion of the sequence space, even if long-lived), so is difficult to consider

Touch

Expires January 8, 2005

[Page 11]

where smaller receive buffers would not sufficiently address the immediate problem.

3.1.6 Other Protocols

Segment authentication has been addressed at the transport layer in other protocols. Both SCTP and DCCP* include cookies for connection establishment and uses them to authenticate a variety of other control messages [14][23]. The inclusion of such mechanism at the transport protocol, although emerging as standard practice, unnecessarily complicates the design and implementation of new protocols. As new attacks are discovered (SYN floods, RSTs, etc.), each protocol must be modified individually to compensate. A network solution may be more appropriate and efficient.

*[AUTH - DCCP may be removing cookies from the spec for the redundancies discussed above, because the use of cookies at the transport layer primarily supports dynamic multihoming (a design goal of SCTP, but not DCCP) rather than security.]

3.2 Network Layer (IP)

TCP is susceptible to RSTs, but also to other spoofing and man-in-the-middle attacks, including SYN attacks. Other transport protocols, such as UDP and RTP are equally susceptible. Although emerging transport protocols attempt to defeat such attacks at the transport layer, it is clear that such attacks are fundamentally a network layer issue. The packet is coming from an endpoint who is spoofing another endpoint, either upstream or somewhere else in the Internet. IPsec was designed specifically to establish and enforce authentication of a packet's source and contents, which most directly, explicitly, and completely addresses this security vulnerability.

The larger problem with IPsec is that of CA key distribution and use. IPsec is often cumbersome, and has only recently been supported in many end-system operating systems. More importantly, it relies on signed X.509 certificates to establish and exchange keying information (e.g., via IKE). These present challenges when using IPsec to secure traffic to a well-known server, whose clients may not support IPsec or may not have registered with a previously-known certificate authority (CA).

Touch

Expires January 8, 2005

[Page 12]

4. Issues

There are a number of existing and proposed solutions addressing the vulnerability of transport protocols in general, and TCP in specific, to off-path third-party spoofing attacks. As shown, these operate at the transport or network layer. These solutions are not a sufficient long-term strategy to dealing with such attacks, however. Transport solutions require pervasive modification of every transport protocol and address the problem of packet origin identification at the wrong layer. Current network solutions are computationally intensive and require pervasive registration of certificate authorities with every possible endpoint. Neither application-layer nor link-layer solutions suffice to protect either the network or transport layers. This section explains these observations in detail.

4.1 Transport Layer (e.g., TCP)

Transport solutions rely on shared cookies to authenticate segments, including data, transport header, and even pseudo-header (e.g., fixed portions of the outer IP header in TCP). Because the Internet relies on stateless network protocols, it makes sense to rely on state establishment and maintenance available in some transport layers not only for the connection but for authentication state. Three-way handshakes and heartbeats can be used to negotiate authentication state in conjunction with connection parameters, which can be stored with connection state easily.

As noted earlier, transport layer solutions require pervasive modification of all transport protocols to include authentication. Not all transport layers support negotiated endpoint state (e.g., UDP), and legacy protocols are notoriously hard to safely augment (e.g., TCP). Not all authentication solutions are created equal either, and relying on a variety of transport solutions exposes end-systems to increased potential for incorrectly specified or implemented solutions. Transport authentication has often been developed piece-wise, in response to specific attacks, e.g., SYN cookies and RST window attenuation [[15](#)][[22](#)].

Transport layer solutions are not only per-protocol, but often per-connection. Each connection needs to negotiate and maintain authentication state separately. Overhead is not amortized over multiple connections - this includes overheads in packet exchanges, design complexity, and implementation complexity. Finally, because the authentication happens later in packet processing than is required, additional endpoint resources are consumed needlessly, e.g., in demultiplexing received packets, indexing connection identifiers, etc.

Touch

Expires January 8, 2005

[Page 13]

4.2 Network Layer (IP)

A network layer solution avoids the hazards of multiple transport variants, using a single shared endpoint authentication mechanism early in receiver packet processing to discard unauthenticated packets quickly. Network solutions protect all transport protocols, including both legacy and emerging protocols, and reduces the complexity of these protocols as well. A shared solution also reduces protocol overhead, and decouples the management (and refreshing) of authentication state from that of individual transport connections. Finally, a network layer solution protects not only the transport layer but the network layer as well, e.g., from ICMP, IGMP, etc. spoofing attacks.

The ubiquitous protocol for network layer authentication is IPsec [16][24]. IPsec specifies the overall architecture, including header authentication (AH) [17][25] and encapsulation (ESP) modes [18]. AH authenticates both the IP header and IP data, whereas ESP authenticates only the IP data (e.g., transport header and payload). AH is deprecated since ESP is more efficient and the SPI includes sufficient information to verify the IP header anyway. These two modes describe the security applied to individual packets within the IPsec system; key exchange and management is performed either out-of-band (via pre-shared keys) or by an automated key exchange protocol IKE [19][26].

IPsec already provides authentication of an IP header and its data contents sufficient to defeat both man-in-the-middle and off-path third-party spoofing attacks. IKE can configure authentication between two endpoints on a per-endpoint, per-protocol, or per-connection basis, as desired. IKE also can perform automatic periodic re-keying, further defeating crypto-analysis based on snooping (clandestine data collection). The use of IPsec is already commonly strongly recommended for protected infrastructure.

IPsec does not suffice for many uses of BGP, however. It is computationally intensive both in key management and individual packet authentication [10]. As importantly, IKE is not anonymous; keys can be exchanged between parties only if they trust each others' X.509 certificates. These certificates provide identification (the other party knows who you are) only where the certificates themselves are signed by certificate authorities (CAs) that both parties already trust. To a large extent, the CAs themselves are the pre-shared keys which help IKE establish security association keys, which are then used in the authentication algorithms.

IPsec, although widely available both in commercial routers and commodity end-systems, is not often utilized except between parties

Touch

Expires January 8, 2005

[Page 14]

that already have a preexisting relationship (employee/employer, between two ISPs, etc.) Servers to anonymous clients (e.g., customer/business) or more open services (e.g., BGP, where routers may large numbers of peers) are unmanageable, due to the breadth and flux of CAs. New endpoints cannot establish IPsec associations with such servers unless their certificate is signed by a CA already trusted by the server. Different servers - even within the same overall system (e.g., BGP) - often cannot or will not trust overlapping subsets of CAs in general.

4.3 Application Layer

There are a number of application layer authentication mechanisms, often implicit within end-to-end encryption. Application-layer security (e.g., TLS, SSH, or MD5 checksums within a BGP stream) provides the ultimate protection of application data from all intermediaries, including network routers as well as exposure at other layers in the end-systems. It is the only way to protect the application data, ultimately.

Application authentication cannot protect either the network or transport protocols from spoofing attacks, however. Spoofed packets interfere with network processing or reset transport connections before the application ever gets to check the data. Authentication needs to winnow these packets and drop them before they interfere at these lower layers.

4.4 Shim Transport/Application Layer

Security can also be provided over the transport layer but below the application layer, in a kind of 'shim' protocol, such as SSL or TLS. These protocols provide data protection for a variety of applications over a single, legacy transport protocol, such as SSL/TCP for HTTPS. Unfortunately, like application authentication, they do not protect the transport layer against spoofing attacks.

4.5 Link Layer

Link layer security operates separately on each hop of an Internet. Such security can be critical in protecting link resources, such as bandwidth and link management protocols. Protection at this layer cannot suffice for network or transport layers, because it cannot authenticate the endpoint source of a packet. Link authentication ensures only the source of the current hop where it is examined.

4.6 Need for Alternate Security Levels

The issues raised in this section suggest the need for alternate

Touch

Expires January 8, 2005

[Page 15]

security levels. While it is already widely recognized that security needs to occur simultaneously at many protocol layers, the need for a variety of strengths at a single layer. IPsec already supports a variety of algorithms (MD5, SHA, etc. for authentication), but always assumes that:

1. the entire body of the packet is secured
2. security associations are established only where identity is authenticated by a known certificate authority or other pre-shared key
3. both man-in-the-middle and off-path third-party spoofing attacks must be defeated

These assumptions are prohibitive, especially in many cases of spoofing attacks. For spoofing, the primary issue is whether packets are coming from the same party the server can reach. Only the IP header is fundamentally in question, so securing the entire packet (1) is computational overkill. It is sufficient to authenticate the other party as "a party you have exchanged packets with", rather than establishing their trusted identity ("Bill" vs. "Bob") as in (2). Finally, many cookie systems use clear-text (unencrypted), fixed cookie values, providing reasonable ($1 \text{ in } 2^{\{\text{cookie-size}\}}$) protection against off-path third-party spoofs, but not addressing man-in-the-middle at all.

Touch

Expires January 8, 2005

[Page 16]

5. The Need for High-Performance Anonymous Security

Internet servers need a form of anonymous security, to protect established connections from spoofing attacks, but without the heavyweight infrastructure typically assumed in conventional security architectures. Such security would allow servers to interact with clients without a-priori shared keys or a key signing hierarchy. It would be useful to also support high-performance variants, especially where only a nonce is sufficient to avoid off-path attacks. Both these properties would address the concerns about deployability of existing IPsec and TCP/MD5 solutions. There are a few different ways to establish anonymous security; the details of these approaches are discussed in detail elsewhere, and are briefly summarized here [4].

5.1 Anonymous Keying

It would be useful to allow IPsec where CAs are not pre-shared, or TCP/MD5 and/or IPsec shared-key mode without a-prior shared keys. The former is a kind of anonymous IKE, where a key is negotiated and exchanged but without requiring pre-agreed CAs. This is already supported by IKE in shared-secret mode, where the shared secret is open, i.e., a public string such as "password". IKE allows such a public shared secret to be used to negotiate a pairwise private secret via the Diffie-Hellman exchange.

TCP/MD5 would need to be augmented to support anonymous automatic keying, since it currently assumes only manual keying. One solution would be to use the TCP connection's Initial Sequence Number (ISN) as the key, since it is visible only during connection establishment. Another solution is to add a Diffie-Hellman exchange to the connection establishment phase, inside the TCP/MD5 option.

Note that the ISN or key established by Diffie-Hellman both represent shared, per-connection information that cannot be obtained after connection establishment or keying, respectively. As such, either can be used as the key itself if only off-path attacks are of concern.

5.2 Alternatives for Performance

An independent impediment to deploying IPsec or TCP/MD5 is performance; in both cases, there is a substantial impact on the throughput of individual connections, as well as the computational load of the endpoints. There are variants of IPsec and TCP/MD5 which can address these performance issues, also discussed in further detail elsewhere [4]. They involve protecting a subset of the packet (focusing on the headers), and/or relaxing the authentication algorithms.

Touch

Expires January 8, 2005

[Page 17]

Conventional authentication involves ESP or ESP-like processing. The entire IP payload is authenticated using existing algorithms, e.g., MD5, SHA, etc. This provides full protection of both all headers and data against all third-party spoofing, as well as tampering. Unfortunately, this alternative is as computationally expensive as other forms of authentication, e.g., TCP MD5, since similar algorithms are used over the bulk of the packet [10][11].

Header-only authentication uses truncated ESP processing over only the fixed portions of the IP (in IPsec) or IP/TCP (in TCP/MD5) header, possibly with some short prefix of the payload, for padding and to avoid small-block hash problems. This mode protects headers from all third-party spoofing and tampering. It does not protect the payload (outside that optionally used to pad the hash) from tampering. Since the IP header is usually sufficiently smaller than a typical hash block (e.g., 64 bytes for MD5), part or all of the TCP header would typically be protected for either IPsec or TCP/MD5. This then completely protects TCP against RST attacks. Header-only mode does not protect from man-in-the-middle replay of signed headers with alternate payloads. For packets substantially larger than the hash block size (again, 64 bytes for MD5), header-only mode can represent a significant computational savings over full-packet authentication.

Both header-only and nonce mode authentication require modifications to IPsec or TCP/MD5. These modes would be negotiated during IKE for IPsec or at connection establishment for TCP/MD5.

Nonce authentication ignores the context of individual packets completely when computing the authentication signature. The signature may be fixed or vary, but is independent of the packet header and contents. The signature is effectively a cookie; it may even be the case that the SPI itself (for IKE) or ISN (for TCP/MD5) are sufficient for this purpose, though we do not recommend either (SPI values are not chosen pseudo-randomly, and ISNs are not always chosen pseudorandomly). The cookie may be the keys established during IKE or the augmented auto-key version of TCP/MD5. The benefit of nonce mode is performance; only header modification is required, avoiding all computational overhead. Nonce mode protects only from off-path third-party spoofing.

Conventional IPsec and TCP/MD5 and header-only mode both protect TCP from all third-party spoofing attacks. Nonce-only mode TCP protects only against off-path third-party attacks, but has much lower computational cost.

In summary, there are two components needed to enable more widespread use of existing IPsec or TCP/MD5 spoofing protection. Automatic

Touch

Expires January 8, 2005

[Page 18]

anonymous keying is already supported by IPsec using IKE with a public shared secret; similar keying can be added to TCP/MD5 by adding a Diffie-Hellman exchange therein, or even by just using the TCP ISN. High-performance variants of both schemes are possible.

6. Security Considerations

This entire document focuses on increasing the security of transport protocols and their resistance to spoofing attacks. Security is addressed throughout.

This document describes a number of techniques for defeating spoofing attacks. Those relying on clear-text cookies, either explicit or implicit (e.g., window sequence attenuation) do not protect from man-in-the-middle spoofing attacks, since valid values can be learned from prior traffic. Those relying on true authentication algorithms are stronger, protecting even from man-in-the-middle, because the authentication hash in a single packet approaches the behavior of "one time" cookies.

Security at various levels of the protocol stack are addressed. Spoofing attacks are fundamentally identity masquerading, so we believe the most appropriate solutions defeat these at the network layer, where end-to-end identity lies. Some transport protocols subsume endpoint identity information from the network layer (e.g., TCP pseudo-headers), while others establish per-connection identity based on exchanged nonces (e.g., SCTP). It is reasonable, if not recommended, to address security at all layers of the protocol stack. We believe that the current attacks are most directly addressed at the network layer.

Some new solutions discussed weaken overall security compared to existing solutions. However, a stronger solution which is not deployed, due to its complexity or performance, is equivalent to no security at all. Providing a more easily deployed anonymous variant of IKE or TCP/MD5 together with authentication that can be adjusted to trade strength for performance may constitute an overall benefit toward the increased deployment of security solutions.

Touch

Expires January 8, 2005

[Page 20]

7. Conclusions

This document describes the details of the recent BGP spoofing attacks involving spurious RSTs used to shutdown TCP connections. It summarizes and discusses a variety of current and proposed solutions at various protocol layers.

8. Acknowledgments

This document was inspired by discussions on the <http://www.ietf.org/html.charters/tcpm-charter.html> about the recent spoofed RST attacks on BGP routers, including R. Stewart's draft [13][22]. The analysis of the attack issues, alternate solutions, and the anonymous security proposed solutions were the result of discussions on that list as well as with USC/ISI's T. Faber, A. Falk, G. Finn, and Y. Wang. Ran Atkinson suggested the UDP variant of TCP/MD5, and Paul Goyette suggested applying headerANON-style processing to TCP/MD5 and using the ISN to seed TCP/MD5. Other improvements are due to the input of various members of the IETF's TCPM WG.

Touch

Expires January 8, 2005

[Page 22]

9. References

9.1 Normative References

- [1] "Technical Cyber Security Alert TA04-111A: Vulnerabilities in TCP -- <http://www.us-cert.gov/cas/techalerts/TA04-111A.html>", , April 20 2004.
- [2] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", [RFC 2385](#), August 1998.
- [3] Poon, K., "Use of TCP timestamp option to defend against blind spoofing attack", (work in progress) , June 2004.
- [4] Touch, J., "ANONsec: Anonymous Security to Defend Against Spoofing Attacks", (work in progress) , July 2004.
- [5] Braden, B., "TIME-WAIT Assassination Hazards in TCP", [RFC 1337](#), May 1992.
- [6] Faber, T., Touch, J. and W. Yue, "The TIME-WAIT state in TCP and Its Effect on Busy Servers", Proc. Infocom 1999 pp. 1573-1583, March 1999.
- [7] Floyd, S., "Inappropriate TCP Resets Considered Harmful", [BCP 60](#), [RFC 3360](#), August 2002.
- [8] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), September 1981.
- [9] Jacobson, V., Braden, B. and D. Borman, "TCP Extensions for High Performance", [RFC 1323](#), May 1992.
- [10] Touch, J., "Report on MD5 Performance", [RFC 1810](#), June 1995.
- [11] Touch, J., "Performance Analysis of MD5", Proc. Sigcomm 1995 77-86., March 1999.
- [12] O'Malley, S. and L. Peterson, "TCP Extensions Considered Harmful", [RFC 1263](#), October 1991.
- [13] "IETF TCPM Working Group and mailing list - <http://www.ietf.org/html.charters/tcpm-charter.html>", .
- [14] Stewart, R., Xie, Q., Morneault, K., Sharp, C., Schwarzbauer, H., Taylor, T., Rytina, I., Kalla, M., Zhang, L. and V. Paxson, "Stream Control Transmission Protocol", [RFC 2960](#), October 2000.

Touch

Expires January 8, 2005

[Page 23]

- [15] Bernstein, D., "SYN cookies -- <http://cr.yp.to/syncookies.html>", , 1997.
- [16] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", [RFC 2401](#), November 1998.
- [17] Kent, S. and R. Atkinson, "IP Authentication Header", [RFC 2402](#), November 1998.
- [18] Kent, S. and R. Atkinson, "IP Encapsulating Security Payload (ESP)", [RFC 2406](#), November 1998.
- [19] Harkins, D. and D. Carrel, "The Internet Key Exchange (IKE)", [RFC 2409](#), November 1998.
- [20] Glenn, R. and S. Kent, "The NULL Encryption Algorithm and Its Use With IPsec", [RFC 2410](#), November 1998.
- [21] Maughan, D., Schneider, M. and M. Schertler, "Internet Security Association and Key Management Protocol (ISAKMP)", [RFC 2408](#), November 1998.

[9.2](#) Informative References

- [22] Stewart, R., "Transmission Control Protocol security considerations", [draft-ietf-tcpm-tcpsecure-01](#) (work in progress), June 2004.
- [23] Kohler, E., "Datagram Congestion Control Protocol (DCCP)", [draft-ietf-dccp-spec-06](#) (work in progress), February 2004.
- [24] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", [draft-ietf-ipsec-rfc2401bis-02](#) (work in progress), April 2004.
- [25] Kent, S., "IP Authentication Header", [draft-ietf-ipsec-rfc2402bis-07](#) (work in progress), March 2004.
- [26] Kaufman, C., "Internet Key Exchange (IKEv2) Protocol", [draft-ietf-ipsec-ikev2-14](#) (work in progress), June 2004.

Touch

Expires January 8, 2005

[Page 24]

Author's Address

Joe Touch
USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292-6695
U.S.A.

Phone: +1 (310) 448-9151
Fax: +1 (310) 448-9300
EMail: touch@isi.edu
URI: <http://www.isi.edu/touch>

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.

Touch

Expires January 8, 2005

[Page 26]