Network Working Group
Internet Draft

Steven M. Bellovin
AT&T Labs Research

Expiration Date: April 2004

October 2003

Guidelines for Mandating the Use of IPsec

draft-bellovin-useipsec-02.txt

Status of this Memo

Abstract

The Security Considerations sections of many Internet Drafts say, in effect, "just use IPsec". While this is sometimes correct, more often it will leave users without real, interoperable security mechanisms. This memo offers some guidance on when IPsec should and should not be specified.

1. Introduction

The Security Considerations sections of many Internet Drafts say, in effect, "just use IPsec". While this is sometimes correct, more often it will leave users without real, interoperable security mechanisms. IPsec is often unavailable in the likely endpoints. Even if it is available, it may not provide the proper granularity of protection. Finally, if it is available and appropriate, the document mandating it needs to specify just how it is to be used.

Recall that the goal is realistic, interoperable security. Specifying, as the only security mechanism, a configuration which is unavailable to -- and hence unusable by -- a majority of the user community is tantamount to saying "turn off security".

For further guidance on security considerations (including discussion of IPsec), see [RFC3552].

NOTE: Many of the arguments below relate to the capabilities of current implementations of IPsec. These may change over time; this advice based on the knowledge available to the IETF at publication time.

2. WARNING

The design of security protocols is a subtle and difficult art. The cautions here about specifying use of IPsec should NOT be taken to mean that that you should invent your own new security protocol for each new application. If IPsec is a bad choice, use another standardized, well-understood security protocol. Don't roll your own.

3. The Pieces of IPsec

IPsec is composed of a number of different pieces. These can be used to provide confidentiality, integrity, and replay protection; though some these can be configured manually, in general a key management component is used. Additionally, the decision about whether and how to use IPsec is controlled by a policy database of some sort.

3.1. AH and ESP

The Authentication Header (AH) [RFC2402] and the Encapsulating Security Protocol (ESP) [RFC2406] are the over-the-wire security protocols. Both provide (optional) replay protection. ESP typically is used to provide confidentiality (encryption), integrity and authentication for traffic. ESP also can provide integrity and authentication without confidentiality, which makes it a good alternative to AH in most cases where confidentiality is not a required or desired service. Finally, ESP can be used to provide confidentiality alone, although this is not recommended [Bell96].

The difference in integrity protection offered by AH is that AH protects portions of the preceding IP header, including the source and destination address. However, when ESP is used in tunnel mode (see section 4.2), and if integrity/authentication is enabled, the IP header seen by the source and destination hosts is completely protected anyway.

AH can also protect those IP options that need to be seen by intermediate routers, but must be intact and authentic when delivered to the receiving system. At this time, use (and existence) of such IP options is extremely rare.

If an application requires such protection, and if the information to be protected cannot be inferred from the key management process, AH must be used. (ESP is generally regarded as easier to implement; however, virtually all IPsec packages support both.) If confidentiality is required, ESP must be used. It is possible to use AH in conjunction with ESP, but this combination is rarely required.

3.2. Transport and Tunnel Mode

AH and ESP can both be used in either transport mode or tunnel mode.

In tunnel mode, the IPsec header is followed by an inner IP header; this is the normal usage for Virtual Private Networks (VPN), and it is generally required whenever either end of the IPsec-protected path is not the ultimate IP destination, e.g., when IPsec is implemented in a firewall, router, etc. Transport mode is preferred for point-to-point communication, though tunnel mode can be used for this purpose.

3.3. Key Management

Any cryptographic system requires key management. IPsec provides for both manual and automatic key management schemes. Manual key management is easy; however, it doesn't scale very well. Also, IPsec's replay protection mechanisms are not available if manual key management is used.

One automated key exchange mechanism is available, Internet Key Exchange (IKE) [RFC2409]. A new, simpler version of IKE is currently being designed. A second mechanism, Kerberized Internet Negotiation of Keys (KINK) [KINK], is being defined. It, of course, uses Kerberos, and is suitable if and only if a Kerberos infrastructure is available.

IKE can use public key certificates or shared secrets. Shared secret authentication is simpler, but doesn't scale as well, since each endpoint must share a unique secret with every peer with which it can communicate, in order to uniquely authenticate these peers.

IKE also has public key variants. In most real-world situations, these use locally-issued certificates. That is, the administrator of the system or network concerned will issue certificates to all authorized users; these certificates are useful only for IPsec.

It is sometimes possible to use certificates from an existing PKI with IKE. In practice, this is rare. Furthermore, there not only is no global PKI for the Internet, there probably never will be one. Designing a structure which assumes such a PKI is a mistake. In

particular, assuming that an arbitrary node will have an "authentic" certificate, issued by a mutually trusted third party and vouching for that node's identity, is wrong. Again, such a PKI does not and probably will not exist. Public key IKE is generally a good idea, but almost always with locally-issued certificates.

Note that public key schemes require a substantial amount of computation. Protocol designers should consider whether or not such computations are feasible on devices of interest to their clientele. Using certificates roughly doubles the number of large exponentiations that must be performed, compared with shared secret versions of IKE.

Today, even low-powered devices can generally perform enough computation to set up a limited number of security associations; concentration points, such as firewalls or VoIP servers, may require hardware assists, especially if many peers are expected to create security associations at about the same time.

3.4. Applications Program Interface (API)

It is, in some sense, a misnomer to speak of the API as a part of IPsec, since that piece is missing on many systems. To the extent that it does exist, it isn't standardized. The problem is simple: it is difficult or impossible to request IPsec protection, or to tell if was used for given inbound packets or connections.

Router- or firewall-based IPsec implementations pose even greater problems, since there is no standardized over-the-wire protocol for communicating this information from outboard encryptors to hosts.

4. Availability of IPsec in Target Devices

Although IPsec is now widely implemented, and is available for current releases of most host operating systems, it is less available for embedded systems. Few hubs, network address translators, etc., implement it, especially at the low end. It is generally inappropriate to rely on IPsec when many of the endpoints are in this category.

Even for host-to-host use, IPsec availability (and experience, and ease of use) has generally been for VPNs. Hosts that support IPsec for VPN use do not always support it on a point-to-point basis, especially via a stable, well-defined API ur user interface.

Finally, few implementations support multiple layers of IPsec. If a telecommuter is using IPsec in VPN mode to access an organizational network, he or she may not be able to employ a second level of IPsec to protect an application connection to a host within the organization. (We note that such support is, in fact, mandated by Case 4 of Section 4.5 of [RFC2041]. Nevertheless, it is not widely available.) The likelihood of such deployment scenarios should be taken into account when deciding whether or not to mandate IPsec.

5. Endpoints

[RFC2401] describes many different forms of endpoint identifier. These include source addresses (both IPv4 and IPv6), host names (possibly as embedded in X.500 certificates), and user IDs (again, possibly as embedded in a certificate). Not all forms of identifier are available on all implementations; in particular, user-granularity identification is not common. This is especially a concern for multi-users systems, where it may not be possible to use different certificates to distinguish between traffic from two different users.

Again, we note that the ability to provide fine-grained protection, such as keying each connection separately, and with per-user credentials, was one of the original design goals of IPsec. Nevertheless, only a few platforms support it. Indeed, some

implementations do not even support using port numbers when deciding whether or not to apply IPsec protection.

6. Selectors and the SPD

Section 4.4 of [RFC2401] describes the Security Policy Database (SPD) and "selectors" used to decide what traffic should be protected by IPsec. Choices include source and destination addresses (or address ranges), protocol numbers (i.e., 6 for TCP and 17 for UDP), and port numbers for TCP and UDP. Protocols whose protection requirements cannot be described in such terms are poorer candidates for IPsecp in particular, it becomes impossible to apply protection at any finer grain than "destination host". Thus, traffic embedded in an L2TP [RFC2661] session cannot be protected selectively by IPsec above the L2TP layer, because IPsec has no selectors defined that let it peer into the L2TP packet to find the TCP port numbers. Similarly, SCTP [RFC2960] did not exist when [RFC2401] was written; thus, protecting individual SCTP applications on the basis of port number could not be done until a new document is written that defines new selectors for IPsec.

The granularity of protection available may have side-effects. If certain traffic between a pair of machines is protected by IPsec, does the implementation permit other traffic to be unprotected, or protected by different policies? Alternatively, if the implementation is such that it is only capable of protecting all traffic or none, does the device have sufficient CPU capacity to encrypt everything? Note that some low-end devices may have limited secure storage capacity for keys, etc.

Implementation issues are also a concern here. As before, too many vendors have not implemented the full specifications; too many IPsec

implementations are not capable of using port numbers in their selectors. Protection of traffic between two hosts is thus on an all or nothing basis when these non-compliant implementations are employed.

7. Broadcast and Multicast

Although the designers of IPsec tried to leave room for protection of multicast traffic, the design has not yet been completed. There is, as yet, no key management for the general case. Worse yet, an important component of over-the-wire IPsec -- replay protection -- is very difficult to implement in a multi-sender situation. IPsec is thus inappropriate for such protocols unless and until suitable key management and replay protection mechanisms are defined and available in the target domain. (Single-sender multicast can be supported, if suitable key management is available; the MSEC working group is developing such a protocol.)

8. Mandating IPsec

Despite all of the caveats given above, it may still be appropriate to use IPsec in particular situations. The range of choices make it mandatory to define precisely how IPsec is to be used. Authors of RFCs that rely on IPsec must specify the following:

(a) What selectors the initiator of the conversation (the client, in client-server architectures) should use? In particular, what addresses, port numbers, etc., are to be used?

(b) What IPsec protocol is to be used: AH or ESP? What mode is to be employed: transport mode or tunnel mode?

(c) What form of key management is appropriate?

(d) What security policy database entry types should be used by the responder (i.e., the server) when deciding whether or not to accept the IPsec connection request.

(e) What form of identification and authentication should be used.

(f) Which of the many variants of IKE must be supported.

(g) Is suitable IPsec support available in likely configurations of the products that would have to employ IPsec?

9. Example

Suppose that the designers of the Border Gateway Protocl (BGP) [RFC1771] wished to use IPsec for security, rather than the mechanism described in [2385]. Does it meet these criteria? (Note that the deeper security issues raised by BGP are not addressed by IPsec or any other transmission security mechanism. See [Kent00a] and [Kent00b] for more details.)

Selectors
: The issue of selectors is easy. BGP already runs between manully-configured pairs of hosts on TCP port 179. The appropriate selector would be the pair of BGP speakers, for that port only. Note that the router's "loopback address" is almost certainly the address to use.

Mode
: Clearly, transport mode is the proper choice. The information being communicated is generally not confidential, so encryption need not be used. Either AH or ESP can be used; if ESP is used, the sender's IP address would need to be checked against the IP address asserted in the key management exchange. (This check is mandated by [RFC2401].)

Key Management
: To permit replay detection, an automated key management system should be used, most likely IKE.

Security Policy
: Connections should be accepted only from the designated peer.

Authentication
: Given the number of BGP-speaking routers used internally by large ISPs, it is likely that shared key mechanisms are inadequate. Consequently, certificate-based IKE must be supported. However, shared secret mode is reasonable on peering links, or (perhaps) on links between ISPs and customers. Whatever scheme is used, it must tie back to a source IP address in some fashion, since other BGP policies are expressed in terms of the peer's IP address.

Availability
: For this scenario, availability is the crucial question. Do likely BGP speakers -- both backbone routers and access routers -- support the profile of IPsec described above? Will use of IPsec, with its attendant expensive cryptographic operations, raise the issue of new denial of service attacks?

The working group and the IESG must make these determinations before deciding to use IPsec to protect BGP.

10. Security Considerations

IPsec provides transmission security and simple access control only. There are many other dimensions to protocol security that are beyond the scope of this memo. Within its scope, the security of any resulting protocol depends heavily on the accuracy of the analysis that resulted in a decision to use IPsec.

11. Acknowledgments

Ran Atkinson, Barbara Fraser, Paul Hoffman, Stephen Kent, Eric Fleischman, and others have made many useful suggestions.

12. References

[Bell96] "Problem Areas for the IP Security Protocols", S.M. Bellovin, Proc. Sixth Usenix Security Symposium, 1996, pp. 205-214.

[Kent00a] "Secure Border Gateway Protocol (Secure-BGP)", S. Kent, C. Lynn, and K. Seo, IEEE Journal on Selected Areas in Communications 18:4, April 2000, pp. 582-592.

[Kent00b] "Secure Border Gateway Protocol (S-BGP) -- Real World Performance and Deployment Issues", S. Kent, C. Lynn, J. Mikkelson, and K. Seo, Proc. Network and Distributed System Security Symposium, February 2000.

[KINK] "Kerberized Internet Negotiation of Keys (KINK)", M. Thomas and J. Vilhuber, work in progress, 2002.

[RFC1771] "A Border Gateway Protocol 4 (BGP-4)", Y. Rekhter and T. Li. RFC 1771, March 1995.

[RFC2401] "Security Architecture for the Internet Protocol", S. Kent and R. Atkinson, RFC 2401, November 1998.

[RFC2402] "IP Authentication Header", S. Kent and R. Atkinson, RFC 2402, November 1998.

Atkinson, RFC 2406, November 1998.

[RFC2409] "The Internet Key Exchange (IKE)", D. Harkins and D. Carrel. RFC 2409, November 1998.

[RFC2661] "Layer Two Tunneling Protocol L2TP", W. Townsley, A. Valencia, A. Rubens, G. Pall, G. Zorn, B. Palter. RFC 2661, August 1999.

[RFC2960] "Stream Control Transmission Protocol", R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. Schwarzbauer, T. Taylor, I. Rytina, M. Kalla, L. Zhang, and V. Paxson. RFC 2960, October 2000.

[RFC3552] "Guidelines for Writing RFC Text on Security Considerations", E. Rescorla and B. Korver, work in progress, 2002.

13. Author Information

Steven M. Bellovin
AT&T Labs Research
Shannon Laboratory
180 Park Avenue
Florham Park, NJ 07932
Phone: +1 973-360-8656
email: bellovin@acm.org