R. Monsour, Hi/fn Inc.
M. Sabin, Consultant
A. Shacham, Cisco Systems
S. Anand, Microsoft Corporation
R. Thayer, Sable Technology

Compression in IP Security

<draft-monsour-compr-ipsec-00.txt>

Status of this Memo

This draft is intended provide input to the IP Security working group as it sorts through the debate regarding the incorporation of lossless data compression into the IP Security architecture. Comments about this draft should be submitted to the authors or to the IPSEC mailing list (ipsec@tis.com).

Abstract

This memo discusses the application of lossless compression technology to the IP Security Architecture [IPSecArch]. Over the last few several months, a number of lively debates on this topic have been held on IPSec mailing list. This memo discusses the issues raised, the alternatives available to resolve them and provides a set of recommendations to bring resolution to the issue.

The goals of the draft are to: (1) define the problem solved by the use of compression in the context of IPSec, (2) review the use of compression and security technologies as they have been applied in other layers of the protocol stack, (3) outline a set of

requirements for applying compression to IPSec, (4) discuss alternative methods which can be implemented to meet the requirements, and (4) propose a set of recommendations for consideration by the working group.

Acknowledgments

The authors wish to acknowledge the many contributors to the compression debate on the IPSec mailing list. In addition, the concept of compressing prior to encrypting data is drawn from several prior sources, including Rodney Thayer's draft [Thayer], the ECP/CCP protocols under PPP and the TLS protocol (better known as SSL).

Table of Contents

1. Introduction

Encrypted data is random in nature and not compressible. When an IP datagram is encrypted, compression methods used at lower protocol layers -- e.g., the PPP Compression Control Protocol [RFC1962] -- no longer work. If both compression and encryption are desired, compression must be performed first.

The IPSec working group of the IETF has been debating the topic of incorporating lossless data compression into the IPSec architecture for the past several months. In fact, the initial introduction of the topic goes as far back as 1994, when Jim Hughes of Network Systems presented the idea at the San Jose meeting of the IPSec working group [Dec96WG].

Following that initial presentation, Network Systems, as well as a handful of other companies have implemented proprietary methods for compressing IP datagrams prior to encrypting them. This memo takes that work, and analyzes similar work done in other security protocols, and proposes a negotiable, interoperable method for applying lossless data compression to the IPSec protocol.

The first compression-related drafts were developed in 1996 and three of those drafts were discussed at the December 1996 IPSec working group meeting in San Jose [Thayer][Sabin1][Sabin2].

2. Compression in IPSec - Problem Definition

The widespread use of the internet has resulted in equally widespread use of point-to-point (PPP) connections between hosts as well as between routers. Lossless compression technology has been deployed in the PPP environment in the form of a sub-protocol known as the Compression Control Protocol (CCP). PPP is a connection-oriented protocol. Many lossless compression technologies have the capability to retain state across each "packet" of data that is compressed. Hence, in a connection-oriented protocol such as

PPP, compression can be applied to the continuing stream of "packets". The principal advantage to such a stateful mechanism is a higher compression ratio.

Another important aspect of the CCP is the negotiability of the compression algorithm for each PPP connection.

Today, PPP is the predominant method for end-user hosts to connect to the internet. Compression in CCP provides end users with significant improvements in bandwidth utilization.

In a different environment, today's routers that support connections in the T1 range AND use lossless compression technology, provide great economic value to their users. For example, a corporate user

employing a leased T1 line can save thousands of dollars per month in line charges through the use of compression technology.

The use of encryption technologies at protocol layers higher than the data-link/PPP layer, renders PPP compression ineffective. With the strong demand for secure communications, encryption is actively being deployed at various layers in the protocol stack to meet the security requirements of various environments. The is the core problem which has driven the compression debate in the IPSec working group.

3. Review of Other Protocols Using Compression and/or Security

This section provides an overview of several examples of other protocols and applications involving the use of lossless compression technology. Some of the protocols described use compression in conjunction with encryption.

3.1 Overview

The diagram below depicts the OSI reference model for data communications protocol layers along with various internet protocols and/or products which incorporate the use of compression and/or security capabilities.

These are just a few examples of the technologies being deployed to meet the security and bandwidth requirements of various classes of users and systems.

Protocol Layers	and/or Products	Compression	Security
Application	PGPmail	yes	yes
Presentation	HTTPv1.1*	yes	no
Session	TLS/SSL	yes	yes
Transport	TCP**	no	no
Network	IPSec	no	yes
Data Link	PPP/CCP/ECP	yes	yes
Physical	link encryptors	no	yes
	link compressors	yes	no

*[RFC2068]

**Note: There is discussion within the IETF of taking up work in this area.

3.2 The Encryption Control Protocol (ECP)

At the same time that the Compression Control Protocol was defined, a sister protocol, the Encryption Control Protocol (ECP) [RFC1968], was defined. In the ECP protocol, it is clearly noted that if compression has been negotiated for a connection, compression must be performed prior to encryption. As far as the authors can tell at this time, the ECP has not been widely deployed. However, it should also be noted that the PPP Extensions working group [PPPEXT] is currently working on additional security protocols in the areas of authentication and public key technologies.

3.3 Transport Layer Security (TLS)

TLS [TLS] (formerly/better known as SSL, the Secure Socket Layer) is a security protocol originally defined and implemented by Netscape Communications Corporation. In the initial draft of TLS 1.0, the use of compression is defined as an optional data transformation which can be used prior to encryption for each packet.

In TLS, the selection of compression algorithm is a negotiated property of the session. Since TLS is a connection-oriented protocol, compression state can be maintained from one packet to the next, thereby improving compression ratio.

3.4 Application Layer Security

There are several examples of application layer security. Clearly, any application which has requirements for confidentiality in its data flow can implement encryption technology to meet such a security requirement. One example of this is an email application. A secure email client is capable of encrypting messages sent from one host to another. PGPmail [PGPMan], a security add-on to many popular email client products, is one such example. PGPmail provides the capability to compress mail prior to encrypting it. No doubt, this is due to the realization of its developers that subsequent attempts to compress the data at lower layers in the protocol stack will be rendered ineffective.

4. Does Compression Fall Within the Scope of the Working Group?

There are several issues which surround the question of whether the IPSec working group should directly consider the issue of applying compression to the IPSec protocols.

The IPSec working group charter specifically states:

> A security protocol in the network layer will be developed to provide cryptographic security services that will flexibly support combinations of authentication, integrity, access control, and confidentiality.
>
> The protocol formats for the IP Authentication Header (AH) and IP Encapsulating Security Payload (ESP) will be independent of the cryptographic algorithm. The preliminary goals will specifically pursue host-to-host security followed by subnet-to-subnet and host-to-subnet topologies.
>
> Protocol and cryptographic techniques will also be developed to support the key management requirements of the network layer security. The Internet Key Management Protocol (IKMP) will be specified as an application layer protocol that is independent of the lower layer security protocol.

The problem definition in section 2 describes the "problem" in terms of the affect of IPSec on a lower-layer protocol. While this is indeed true, it is unclear whether the obligation to provide the detailed protocol "fix" to correct the problem lies within the scope of the IPSec working group.

While progress is being made, it is also true that the IPSec specifications have been not been stable and available in a manner to support widespread interoperable implementations (as has been noted by various members of the working group).

The user community has indicated their requirement that compression capabilities be "supported" by IPSec implementations. The specific methods used to provide such support are clearly in the scope of the working group to decide.

There is currently discussion underway (within the IETF) to explore the application of compression in TCP. This raises a question regarding the potential to provide a consistent mechanism for supporting compression across these two closely-related protocols (TCP & IP).

5. Requirements for Applying Compression to IPSec

As noted earlier, the use of encryption at any protocol layer above the data link layer renders PPP compression ineffective. This leads to the need to support the use of compression in the context of IPSec. The question then becomes one of how to provide this support.

An understanding of the problem, the approaches taken in other protocol environments, the emerging discussion of compression support in TCP, and the discussions on the mailing list lead to the definition of the following requirements for the technical

approaches to support compression in IPSec.

5.1 Negotiated Algorithm

Regardless of the protocol for which compression support is provided, a method is required for the two communicating parties to negotiate both the use and the type of compression to be applied to the packets.

5.2 Stateless AND Stateful Compression

Since IP is a connectionless/stateless protocol, it is important that each compressed packet be capable of being decompressed independently of any other packet; i.e., the successful decompression of a packet must not depend on the contents of any other packet(s), nor should it depend on order of receipt of any other packet(s).

The development of a consistent approach to providing compression capabilities to both IPSec and TCP should support the use of stateful compression in the TCP context. TCP's connection orientation can guarantee packet ordering and thus achieve higher compression ratios.

5.3 Compressed Packet Indicator

Once the use of compression is negotiated between two communicating parties, the sender must have the flexibility to determine whether or not to compress each packet. Such flexibility requires a per-packet indication of whether or not the packet is compressed.

5.4 Compatibility with Existing and Emerging Implementations

Changes to the IPSec protocol definitions to support compression must not not break any existing implementation. This means that for an implementation to be compression-capable, it will have to be modified, but it will remain compliant with the IPSec protocols without the addition of compression support. This requirement becomes more desirable with the passage of time, as more and more implementations are developed and deployed.

5.5 Optional Use of Compression

This requirement derives from the previous requirement. If an existing implementation is to be considered compliant with the IPSec protocol, then it MUST NOT be required to provide support for compression.

5.6 Ability to Optimize Compression Across Protocol Layers

The use of compression at several layers in the protocol stack can cause inefficiencies in the processing by sending systems. For example, in an environment where application layer compression is in use, it is desirable to communicate to the lower layer protocols that the data is already compressed and that the lower layers need not attempt to compress (read as "waste cycles"). Thus, support for compression in IPSec shall be provide the capability for "turning compression on or off" through some mechanism of inter-layer communication.

5.7 Anti-Expansion of Payload Data

It is not possible in all instances to pre-determine if a payload has already been compressed at a higher layer. Future protocol changes could support such a pre-determination. Until then, a mechanism for detecting the expansion of data and optionally sending the original uncompressed payload is required.

6. Alternative Approaches to the Use of Compression in IPSec

Two general technical approaches to meet the requirements outlined in previous section have been presented to the working group and subsequently debated on the mailing list. These two approaches are described below along with their relative advantages and disadvantages.

6.1 Layered Architecture Using a Separate Compression Header

This approach involves the use of a separate compression header which would follow the IP header and precede the AH and/or ESP header(s). The header would provide all the fields necessary to support compression of either AH and/or ESP payloads as well as support for compression in TCP.

6.2 Optional Feature of ESP

This approach is based on the fact that it is the encryption of ESP payloads which renders downstream compression techniques ineffective. This approach is limited to only compressing ESP payloads and does not extend naturally to any upper layer protocols.

6.3 Comparison of the Two Approaches

Below is a comparison of the advantages and disadvantages of the two approaches.

	Advantages	Disadvantages
Layered Architecture	not tied to use of encryption; i.e., works with AH payloads as well can be applied to upper layer protocols such as TCP in addition to IPSec routers can more easily determine if a packet is compressed without knowledge of IPSec protocol details	some delay in getting to a standard
Optional Feature of ESP	Easily defined due to its limited scope	doesn't map well for use in upper layer protocols tied to use of security protocols requires routers to know IPSec to avoid compression processing overhead

7. Which Approach Meets the Proposed Requirements?

This section describes how the two approaches described in the previous section meet each of the requirements defined in section 5.

7.1 Negotiated Algorithm

Both approaches meet this requirement. ISAKMP provides a method for algorithm negotiation which can easily be extended to support the negotiation of the use and type of compression. In the TCP context, TCP negotiation would be extended to negotiate the use and type of compression.

7.2 Stateless AND Stateful Compression

IPSec compression MUST be stateless. Both approaches support this concept.

TCP, as a connection-oriented protocol, can support stateful compression and achieve higher compression ratios as a result. The use of the layered architecture approach makes stateful TCP compression possible.

7.3 Compressed Packet Indicator

In the layered architecture approach, the compression header would contain a field to indicate whether or not the packet is compressed.

In the ESP option approach, the most-significant bit of the pad length field has been proposed to serve this function.

7.4 Compatibility with Existing and Emerging Implementations

Since compression is a negotiated option in both approaches, they both meet this requirement. Existing implementations will not negotiate to use compression and will continue to interoperate with new and existing IPSec compliant implementations.

7.6 Ability to Optimize Compression Across Protocol Layers

The layered architecture approach meets this requirement.

The ESP option approach sacrifices the opportunity to develop a single method for compressing across IP and TCP layer protocol.

7.7 Anti-Expansion of Payload Data

Both approaches can meet this requirement.

8. Conclusions

The authors have drawn the following conclusions based on discussions with user community members, analysis of the proposed technical approaches, and the emerging need to broaden the use of compression beyond the IPSec context.

- The need for compression in the IPSec context exists. The effect of IPSec encryption on downstream compression and the demands by members of the user community demonstrate a clear need for supporting compression in an IPSec context.

- The compression topic does not distinctly fall in the

charter of the IPSec working group.

- A layered architecture approach is a superior approach when compared to the ESP option approach to problem of providing compression capabilities in the IPSec context.

9. Recommendations

The authors recommend the following course of action.

- Document the IPSec Architecture to Support Optional Compression

  It is important that the IPSec architecture specification include the notion that payload data of IPSec payloads may optionally be compressed. The draft should note that the specification for providing such compression capabilities will be developed outside of the IPSec working group.

- Develop a Layered Architecture Approach to Compression

  A layered architecture approach, when considered against the ESP option approach has significant advantages which outweigh any potential delays in specification development and subsequent deployment.

- Establish a compression working group

  This group would take on the responsibility for defining a the layered architecture and the separate compression header for use in IPSec as well as other candidate upper layer protocols.

  The rationale for this recommendation is based on several factors, including (1) the user community as well as the vendor community are under pressure to finalize the IPSec specifications, complete development and begin deployment of IPSec-compliant products. As Hugo put in a recent post to the list, "...further delay of ipsec deployment (which is also a form of denial-of-service attack :)", (2) the desire to support compression in a context broader than IPSec alone (i.e., TCP or other candidate protocols), and

10. Security Considerations

This memo discusses the use of lossless compression technology in a security protocol, specifically IPSec. The proposed use of compression within this protocol is not believed to have an effect on the underlying security functionality provide by the protocol; i.e., the use of compression is not known to degrade or alter the nature of the underlying security architecture or the encryption technologies used to implement it.

The use of compression does change the length of ESP payloads, in a manner that depends on the data prior to encryption. Thus, the use of compression may have an effect on the ability of an eavesdropper to glean information by analyzing the length of transmitted packets.

11. References

[IPSecArch] Atkinson, R., "Security Architecture for the Internet Protocol," November 1996.

[Dec96WG] Lambert, P., Minutes of the IPSec Working Group, San Jose December 1994.

[Thayer] Thayer, R., "Compression Header for IPSec", draft-thayer-seccomp-00.txt, May, 1996.

[Sabin1] Sabin, M., Monsour R., "

[Sabin2] Sabin, M., Monsour R., "

[RFC1962] Rand, D., "The PPP Compression Control Protocol (CCP)," RFC-1962, June 1996.

[RFC2068] ????, "Hypertext Transport Protocol version 1.1", January, 1997.

[RFC1968] ????, "The PPP Encryption Control Protocol (ECP)", June 1996.

[PPPEXT] Point-to-Point Protocol Extensions Working Group Charter.

[TLS] T. Dierks, C. Allen, "The TLS Protocol Version 1.0", draft-ietf-tls-protocol-01.txt, March 1997.

[PGPMan] ????, "PGPmail Reference Manual"

[AH] Atkinson, R., "IP Authentication Header"

[ESP] Atkinson, R., "IP Encapsulating Security Payload," June 1996.

[DOI] Piper, D., "IPSec Domain of Interpretation", March 1997.

[Calgary] Text Compression Corpus, University of Calgary, available at ftp://ftp.cpsc.ucalgary.ca/pub/projects/text.compression.corpus.

8. Authors' Addresses

Robert Monsour
Hi/fn Inc.
12636 High Bluff Drive
San Diego, CA 92130
Email: rmonsour@hifn.com

Michael Sabin
883 Mango Avenue
Sunnyvale, CA 94087

Email: mike.sabin@worldnet.att.net

Abraham Shacham
Cisco Systems
101 Cooper Street
Santa Cruz, CA 95060
Email: shacham@cisco.com

Sanjay Anand
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052-6399
Email: sanjayan@microsoft.com

Rodney Thayer
Sable Technology
246 Walnut Street
Newton, MA 02160
Email: rodney@sabletech.com