R. Atkinson
Naval Research Laboratory

IP Authentication Header

Status of this Memo

ABSTRACT

This document describes a mechanism for providing cryptographic authentication for IPv4 and IPv6 datagrams. An Authentication Header (AH) is normally inserted after an IP header and before the other information being authenticated.

1. INTRODUCTION

The Authentication Header is a mechanism for providing strong integrity and authentication for IP datagrams. It might also provide non-repudiation, depending on which cryptographic algorithm is used and how keying is performed. For example, use of an asymmetric digital signature algorithm, such as RSA, could provide non-repudiation.

Confidentiality, and protection from traffic analysis are not provided by the Authentication Header. Users desiring confidentiality should consider using the IP Encapsulating Security Protocol (ESP) either in lieu of or in conjunction with the Authentication Header [Atk95b]. This document assumes the reader has previously read the related IP Security Architecture document which defines the overall security architecture for IP and provides important background information for this specification [Atk95a].

1.1 Overview

The IP Authentication Header seeks to provide security by adding authentication information to an IP datagram. This authentication information is calculated using all of the fields in the IP datagram (including not only the IP Header but also other headers and the user data) which do not change in transit. Fields or options which need to change in transit (e.g., "hop count", "time to live", "ident",

"fragment offset", or "routing pointer") are considered to be zero for the calculation of the authentication data. This provides significantly more security than is currently present in IPv4 and might be sufficient for the needs of many users.

Use of this specification will increase the IP protocol processing costs in participating end systems and will also increase the communications latency. The increased latency is primarily due to the calculation of the authentication data by the sender and the calculation and comparison of the authentication data by the receiver for each IP datagram containing an Authentication Header. The impact will vary with authentication algorithm used and other factors.

In order for the Authentication Header to work properly without changing the entire Internet infrastructure, the authentication data is carried in its own payload. Systems that aren't participating in the authentication MAY ignore the Authentication Data. When used with IPv6, the Authentication Header is normally placed after the Fragmentation and End-to-End headers and before the ESP and transport-layer headers. The information in the other IP headers is used to route the datagram from origin to destination. When used with IPv4, the Authentication Header immediately follows an IPv4 header.

If a symmetric authentication algorithm is used and intermediate authentication is desired, then the nodes performing such intermediate authentication would need to be provided with the appropriate keys. Possession of those keys would permit any one of those systems to forge traffic claiming to be from the legitimate sender to the legitimate receiver or to modify the contents of otherwise legitimate traffic. In some environments such intermediate authentication might be desirable [BCCH94]. If an asymmetric authentication algorithm is used and the routers are aware of the appropriate public keys and authentication algorithm, then the routers possessing the authentication public key could authenticate the traffic being handled without being able to forge or modify otherwise legitimate traffic. Also, Path MTU Discovery MUST be used when intermediate authentication of the Authentication Header is desired and IPv4 is in use because with this method it is not possible to authenticate a fragment of a packet [MD90] [Kno93].

1.2 Requirements Terminology

In this document, the words that are used to define the significance of each particular requirement are usually capitalised. These words are:

- MUST

  This word or the adjective "REQUIRED" means that the item is an absolute requirement of the specification.

- SHOULD

  This word or the adjective "RECOMMENDED" means that there might exist valid reasons in particular circumstances to ignore this item, but the full implications should be understood and the case carefully weighed before taking a different course.

- MAY

  This word or the adjective "OPTIONAL" means that this item is truly optional. One vendor might choose to include the item because a particular marketplace requires it or because it enhances the product, for example; another vendor may omit the same item.

2. KEY MANAGEMENT

Key management is an important part of the IP security architecture. However, it is not integrated with this specification because of a long history in the public literature of subtle flaws in key management algorithms and protocols. The IP Authentication Header tries to decouple the key management mechanisms from the security protocol mechanisms. The only coupling between the key management protocol and the security protocol is with the Security Parameters Index (SPI), which is described in more detail below. This decoupling permits several different key management mechanisms to be used. More importantly, it permits the key management protocol to be changed or corrected without unduly impacting the security protocol implementations.

The key management mechanism is used to negotiate a number of parameters for each "Security Association", including not only the keys but also other information (e.g., the authentication algorithm and mode) used by the communicating parties. The key management mechanism creates and maintains a logical table containing the several parameters for each current security association. An implementation of the IP Authentication Header will need to read that

logical table of security parameters to determine how to process each datagram containing an Authentication Header (e.g., to determine which algorithm/mode and key to use in authentication).

Security Associations are unidirectional. A bidirectional communications session will normally have one Security Association in each direction. For example, when a TCP session exists between two systems A and B, there will normally be one Security Association from A to B and a separate second Security Assocation from B to A. The receiver assigns the SPI value to the the Security Association with that sender. The other parameters of the Security Association are determined in a manner specified by the key management mechanism. Section 4 of this document describes in detail the process of selecting a Security Association for an outgoing packet and identifying the Security Assocation for an incoming packet.

The IP Security Architecture document describes key management in detail. It includes specification of the key management requirements for this protocol, and is incorporated here by reference [Atk95a].

3. AUTHENTICATION HEADER SYNTAX

The Authentication Header (AH) may appear after any other headers which are examined at each hop, and before any other headers which are not examined at an intermediate hop. The IPv4 or IPv6 header immediately preceding the Authentication Header will contain the value 51 in its Next Header (or Protocol) field [STD-2].

Example high-level diagrams of IP datagrams with the Authentication Header follow.

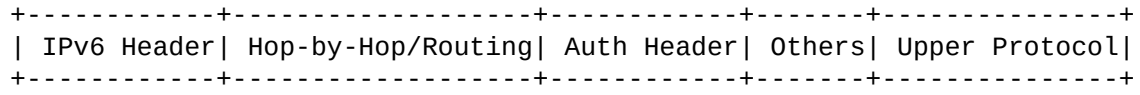

Figure 1: IPv6 Example

When used with IPv6, the Authentication Header normally appears after the IPv6 Hop-by-Hop Header and before the IPv6 Destination Options.

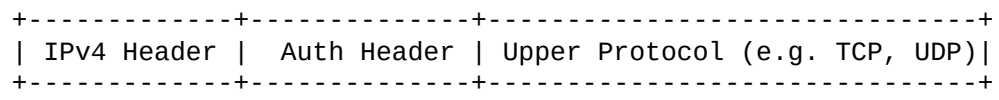

Figure 2: IPv4 Example

When used with IPv4, the Authentication Header normally follows the main IPv4 header.

3.1 Authentication Header Syntax

The authentication data is the output of the authentication algorithm calculated over the the entire IP datagram as described in more detail later in this document. The authentication calculation must treat the Authentication Data field itself and all fields that are normally modified in transit (e.g., TTL or Hop Limit) as if those fields contained all zeros. All other Authentication Header fields are included in the authentication calculation normally.

The IP Authentication Header has the following syntax:

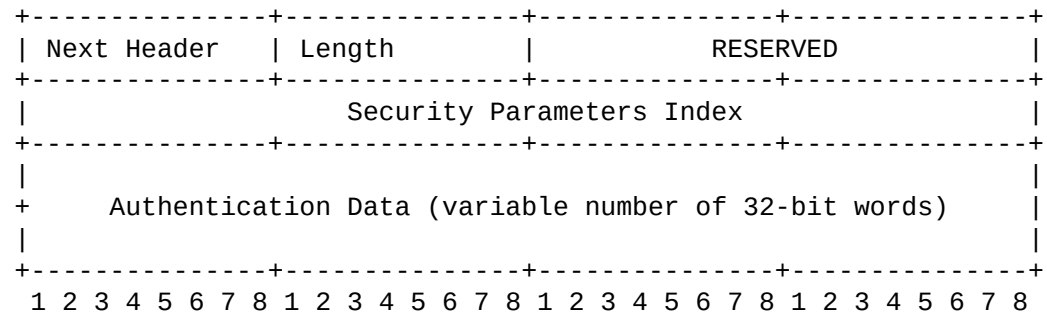

Figure 3: Authentication Header syntax

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.