

RIP-2 MD5 Authentication

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Table of Contents

1 Use of Imperatives	1
2 Introduction	2
3 Implementation Approach	3
3.1 RIP-2 PDU Format	3
3.2 Processing Algorithm	5
3.2.1 Message Generation	6
3.2.2 Message Reception	7
4 Management Procedures	7
4.1 Key Management Requirements	7
4.2 Key Management Procedures	8
4.3 Pathological Cases	9
5 Conformance Requirements	9
6 Acknowledgments	10
7 References	10
8 Security Considerations	11
9 Chairman's Address	11
10 Authors' Addresses	12

1. Use of Imperatives

Throughout this document, the words that are used to define the significance of particular requirements are capitalized. These words are:

MUST

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

MUST NOT

This phrase means that the item is an absolute prohibition of this specification.

SHOULD

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

SHOULD NOT

This phrase means that there may exist valid reasons in particular circumstances when the listed behavior is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behavior described with this label.

MAY

This word or the adjective "OPTIONAL" means that this item is truly optional. One vendor may 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. Introduction

Growth in the Internet has made us aware of the need for improved authentication of routing information. RIP-2 provides for unauthenticated service (as in classical RIP), or password authentication. Both are vulnerable to passive attacks currently widespread in the Internet. Well-understood security issues exist in routing protocols [4]. Clear text passwords, currently specified for use with RIP-2, are no longer considered sufficient [5].

If authentication is disabled, then only simple misconfigurations are detected. Simple passwords transmitted in the clear will further protect against the honest neighbor, but are useless in the general case. By simply capturing information on the wire - straightforward even in a remote environment - a hostile process can learn the password and overcome the network.

We propose that RIP-2 use an authentication algorithm, as was originally proposed for SNMP Version 2, augmented by a sequence number. Keyed MD5 is proposed as the standard authentication algorithm for RIP-2, but the mechanism is intended to be algorithm-independent. While this mechanism is not unbreakable (no known

mechanism is), it provides a greatly enhanced probability that a system being attacked will detect and ignore hostile messages. This is because we transmit the output of an authentication algorithm (e.g., Keyed MD5) rather than the secret RIP-2 Authentication Key. This output is a one-way function of a message and a secret RIP-2 Authentication Key. This RIP-2 Authentication Key is never sent over the network in the clear, thus providing protection against the passive attacks now commonplace in the Internet.

In this way, protection is afforded against forgery or message modification. It is possible to replay a message until the sequence number changes, but the sequence number makes replay in the long term less of an issue. The mechanism does not afford confidentiality, since messages stay in the clear; however, the mechanism is also exportable from most countries, which test a privacy algorithm would fail.

Other relevant rationales for the approach are that Keyed MD5 is being used for OSPF cryptographic authentication, and is therefore present in routers already, as is some form of password management. A similar approach has been standardized for use in IP-layer authentication. [7]

3. Implementation Approach

Implementation requires three issues to be addressed:

- (1) A changed packet format,
- (2) Authentication procedures, and
- (3) Management controls.

3.1. RIP-2 PDU Format

The basic RIP-2 message format provides for an 8 byte header with an array of 20 byte records as its data content. When Keyed MD5 is used, the same header and content are used, except that the 16 byte "authentication key" field is reused to describe a "Keyed Message Digest" trailer. This consists in five fields:


- (1) The "Authentication Type" is Keyed Message Digest Algorithm, indicated by the value 3 (1 and 2 indicate "IP Route" and "Password", respectively).
- (2) A 16 bit offset from the RIP-2 header to the MD5 digest (if no other trailer fields are ever defined, this value equals the RIP-2 Data Length).

- (3) An unsigned 8-bit field that contains the Key Identifier or Key-ID. This identifies the key used to create the Authentication Data for this RIP-2 message. In implementations supporting more than one authentication algorithm, the Key-ID also indicates the authentication algorithm in use for this message. A key is associated with an interface.
- (4) An unsigned 8-bit field that contains the length in octets of the trailing Authentication Data field. The presence of this field permits other algorithms (e.g., Keyed SHA) to be substituted for Keyed MD5 if desired.
- (5) An unsigned 32 bit sequence number. The sequence number MUST be non-decreasing for all messages sent with the same Key ID.


The trailer consists of the Authentication Data, which is the output of the Keyed Message Digest Algorithm. When the Authentication Algorithm is Keyed MD5, the output data is 16 bytes; during digest calculation, this is effectively followed by a pad field and a length field as defined by RFC 1321.

3.2. Processing Algorithm

When the authentication type is "Keyed Message Digest", message processing is changed in message creation and reception.



In memory, the following trailer is appended by the MD5 algorithm and treated as though it were part of the message.



3.2.1. Message Generation

The RIP-2 Packet is created as usual, with these exceptions:

- (1) The UDP checksum need not be calculated, but MAY be set to zero.
- (2) The authentication type field indicates the Keyed Message Digest Algorithm (3).
- (3) The authentication "password" field is reused to store a packet offset to the Authentication Data, a Key Identifier, the Authentication Data Length, and a non-decreasing sequence number.

The value used in the sequence number is arbitrary, but two suggestions are the time of the message's creation or a simple message counter.

The RIP-2 Authentication Key is selected by the sender based on the outgoing interface. Each key has a lifetime associated with it. No key is ever used outside its lifetime. Since the key's algorithm is related to the key itself, stored in the sender and receiver along with it, the Key ID effectively indicates which authentication algorithm is in use if the implementation supports more than one authentication algorithm.

- (1) The RIP-2 header's packet length field indicates the standard RIP-2 portion of the packet.
- (2) The Authentication Data Offset, Key Identifier, and Authentication Data size fields are filled in appropriately.
- (3) The RIP-2 Authentication Key, which is 16 bytes long when the Keyed MD5 algorithm is used, is now appended to the data. For all algorithms, the RIP-2 Authentication Key is never longer than the output of the algorithm in use.
- (4) Trailing pad and length fields are added and the digest calculated using the indicated algorithm. When Keyed MD5 is the algorithm in use, these are calculated per RFC 1321.
- (5) The digest is written over the RIP-2 Authentication Key. When MD5 is used, this digest will be 16 bytes long.

The trailing pad is not actually transmitted, as it is entirely predictable from the message length and algorithm in use.

3.2.2. Message Reception

When the message is received, the process is reversed:

- (1) The digest is set aside,
- (2) The appropriate algorithm and key are determined from the value of the Key Identifier field,
- (3) The RIP-2 Authentication Key is written into the appropriate number (16 when Keyed MD5 is used) of bytes starting at the offset indicated,
- (4) Appropriate padding is added as needed, and
- (5) A new digest calculated using the indicated algorithm.

If the calculated digest does not match the received digest, the message is discarded unprocessed. If the neighbor has been heard from recently enough to have viable routes in the route table and the received sequence number is less than the last one received, the message likewise is discarded unprocessed. When connectivity to the neighbor has been lost, the receiver SHOULD be ready to accept either:

- a message with a sequence number of zero
- a message with a higher sequence number than the last received sequence number.

A router that has forgotten its current sequence number but remembers its key and Key-ID MUST send its first packet with a sequence number of zero. This leaves a small opening for a replay attack. Router vendors are encouraged to provide stable storage for keys, key lifetimes, Key-IDs, and the related sequence numbers.

Acceptable messages are now truncated to RIP-2 message itself and treated normally.

4. Management Procedures

4.1. Key Management Requirements

It is strongly desirable that a hypothetical security breach in one Internet protocol not automatically compromise other Internet protocols. The Authentication Key of this specification SHOULD NOT be stored using protocols or algorithms that have known flaws.

Implementations MUST support the storage of more than one key at the same time, although it is recognized that only one key will normally be active on an interface. They MUST associate a specific lifetime (i.e., date/time first valid and date/time no longer valid) and a key identifier with each key, and MUST support manual key distribution (e.g., the privileged user manually typing in the key, key lifetime, and key identifier on the router console). The lifetime may be infinite. If more than one algorithm is supported, then the implementation MUST require that the algorithm be specified for each key at the time the other key information is entered. Keys that are out of date MAY be deleted at will by the implementation without requiring human intervention. Manual deletion of active keys SHOULD also be supported.

It is likely that the IETF will define a standard key management protocol. It is strongly desirable to use that key management protocol to distribute RIP-2 Authentication Keys among communicating RIP-2 implementations. Such a protocol would provide scalability and significantly reduce the human administrative burden. The Key ID can be used as a hook between RIP-2 and such a future protocol. Key management protocols have a long history of subtle flaws that are often discovered long after the protocol was first described in public. To avoid having to change all RIP-2 implementations should such a flaw be discovered, integrated key management protocol techniques were deliberately omitted from this specification.

4.2. Key Management Procedures

As with all security methods using keys, it is necessary to change the RIP-2 Authentication Key on a regular basis. To maintain routing stability during such changes, implementations MUST be able to store and use more than one RIP-2 Authentication Key on a given interface at the same time.

Each key will have its own Key Identifier, which is stored locally. The combination of the Key Identifier and the interface associated with the message uniquely identifies the Authentication Algorithm and RIP-2 Authentication Key in use.

As noted above in Section 2.2.1, the party creating the RIP-2 message will select a valid key from the set of valid keys for that interface. The receiver will use the Key Identifier and interface to determine which key to use for authentication of the received message. More than one key may be associated with an interface at the same time.

Hence it is possible to have fairly smooth RIP-2 Authentication Key rollovers without losing legitimate RIP-2 messages because the stored key is incorrect and without requiring people to change all the keys at once. To ensure a smooth rollover, each communicating RIP-2 system must be updated with the new key several minutes before the current key will expire and several minutes before the new key lifetime begins. The new key should have a lifetime that starts several minutes before the old key expires. This gives time for each system to learn of the new RIP-2 Authentication Key before that key will be used. It also ensures that the new key will begin being used and the current key will go out of use before the current key's lifetime expires. For the duration of the overlap in key lifetimes, a system may receive messages using either key and authenticate the message. The Key-ID in the received message is used to select the appropriate key for authentication.

4.3. Pathological Cases

Two pathological cases exist which must be handled, which are failures of the network manager. Both of these should be exceedingly rare.

During key switchover, devices may exist which have not yet been successfully configured with the new key. Therefore, routers SHOULD implement (and would be well advised to implement) an algorithm that detects the set of keys being used by its neighbors, and transmits its messages using both the new and old keys until all of the neighbors are using both the new key or the lifetime of the old key expires. Under normal circumstances, this elevated transmission rate will exist for a single update interval.

In the event that the last key associated with an interface expires, it is unacceptable to revert to an unauthenticated condition, and not advisable to disrupt routing. Therefore, the router should send a "last authentication key expiration" notification to the network manager and treat the key as having an infinite lifetime until the lifetime is extended, the key is deleted by network management, or a new key is configured.

5. Conformance Requirements

To conform to this specification, an implementation MUST support all of its aspects. The Keyed MD5 authentication algorithm MUST be implemented by all conforming implementations. MD5 is defined in RFC-1321. A conforming implementation MAY also support other authentication algorithms such as Keyed Secure Hash Algorithm (SHA). Manual key distribution as described above MUST be supported by all conforming implementations. All implementations MUST support the

smooth key rollover described under "Key Change Procedures."

The user documentation provided with the implementation MUST contain clear instructions on how to ensure that smooth key rollover occurs.

Implementations SHOULD support a standard key management protocol for secure distribution of RIP-2 Authentication Keys once such a key management protocol is standardized by the IETF.

6. Acknowledgments

This work was done by the RIP-2 Working Group, of which Gary Malkin is the Chair. This suggestion was originally made by Christian Huitema on behalf of the IAB. Jeff Honig (Cornell) and Dennis Ferguson (ANS) built the first operational prototype, proving out the algorithms. The authors gladly acknowledge significant inputs from each of these sources.

7. References

- [1] Malkin, G., "RIP Version 2 Carrying Additional Information", RFC 1388, January 1993.
- [2] Rivest, R., "The MD5 Message-Digest Algorithm", RFC 1321, April 1992.
- [3] Malkin, G., and F. Baker, "RIP Version 2 MIB Extension", RFC 1389, Xylogics, Inc., Advanced Computer Communications, January 1993.
- [4] S. Bellovin, "Security Problems in the TCP/IP Protocol Suite", ACM Computer Communications Review, Volume 19, Number 2, pp.32-48, April 1989.
- [5] Haller, N., and R. Atkinson, "Internet Authentication Guidelines", RFC 1704, October 1994.
- [6] Braden, R., Clark, D., Crocker, S., and C. Huitema, "Report of IAB Workshop on Security in the Internet Architecture", RFC 1636, June 1994.
- [7] Atkinson, R., "IP Authentication Header", RFC 1826, August 1995.
- [8] Atkinson, R., "IP Encapsulating Security Payload", RFC 1827, August 1995.

8. Security Considerations

This entire memo describes and specifies an authentication mechanism for the RIP-2 routing protocol that is believed to be secure against active and passive attacks. Passive attacks are clearly widespread in the Internet at present. Protection against active attacks is also needed because active attacks are becoming more common.

Users need to understand that the quality of the security provided by this mechanism depends completely on the strength of the implemented authentication algorithms, the strength of the key being used, and the correct implementation of the security mechanism in all communicating RIP-2 implementations. This mechanism also depends on the RIP-2 Authentication Key being kept confidential by all parties. If any of these incorrect or insufficiently secure, then no real security will be provided to the users of this mechanism.

Specifically with respect to the use of SNMP, compromise of SNMP security has the necessary result that the various RIP-2 configuration parameters (e.g. routing table, RIP-2 Authentication Key) manageable via SNMP could be compromised as well. Changing Authentication Keys using non-encrypted SNMP is no more secure than sending passwords in the clear.

Confidentiality is not provided by this mechanism. Recent work in the IETF provides a standard mechanism for IP-layer encryption. [8] That mechanism might be used to provide confidentiality for RIP-2 in the future. Protection against traffic analysis is also not provided. Mechanisms such as bulk link encryption might be used when protection against traffic analysis is required.

The memo is written to address a security consideration in RIP Version 2 that was raised during the IAB's recent security review [6].

9. Chairman's Address

Gary Scott Malkin
Xylogics, Inc.
53 Third Avenue
Burlington, MA 01803

Phone: (617) 272-8140
EMail: gmalkin@Xylogics.COM

10. Authors' Addresses

Fred Baker
cisco Systems
519 Lado Drive
Santa Barbara, California 93111

Phone: (805) 681 0115
Email: fred@cisco.com

Randall Atkinson
cisco Systems
170 West Tasman Drive
San Jose, CA 95134-1706

Phone: (408) 526-6566
EMail: rja@cisco.com

