

Network Working Group
Internet-Draft
Intended Status: Informational
Expires: December 9, 2010

Vishwas Manral
IP Infusion
Manav Bhatia
Alcatel-Lucent
Joel Jaeggli
Checkpoint Software
Russ White
Cisco Systems
June 9, 2010

Issues with existing Cryptographic Protection Methods for Routing Protocols

[draft-ietf-opsec-routing-protocols-crypto-issues-06.txt](#)

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#). This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

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/lid-abstracts.html>

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

This Internet-Draft will expire on December 1, 2010.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

Manral, et. al.

[Page 1]

Internet-Draft

Routing Protocol Protection Issues

June 2010

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Abstract

Routing protocols have over time been extended to use cryptographic mechanisms to validate data being received from a neighboring router to ensure that:

- o It has not been modified in transit.
- o Actually originated from an authorized neighboring router.

The cryptographic mechanisms defined to date and described in this document rely on a digest produced with a hash algorithm applied to the payload encapsulated in the routing protocol packet.

This document outlines some of the limitations of the current mechanism, problems with manual keying of these cryptographic algorithms, and possible vectors for the exploitation of these limitations.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#). [[RFC2119](#)]

Table of Contents

1.	Problem Statement.....	3
1.1	MD5 Pre-Image vs Collision Attacks.....	4

1.2 Concerns about MD5 and the SHA-1 Algorithm.....	4
2. Open Shortest Path First Version 2 (OSPFv2).....	5
2.1 Management Issues with OSPFv2.....	5
2.2 Technical Issues with OSPFv2.....	5
3. Open Shortest Path First Version 3(OSPFv3).....	7
3.1 Management Issues with OSPFv3.....	7
3.2 Technical Issues with OSPFv3.....	8
4. Intermediate System to Intermediate System Routing Protocol (IS-IS).....	9
4.1 Management Issues with IS-IS.....	9
4.2 Technical Issues with IS-IS.....	9

Manral, et. al. [Page 2]

Internet-Draft

Routing Protocol Protection Issues

June 2010

5. Border Gateway Protocol (BGP-4).....	11
5.1 Management Issues with BGP-4.....	12
5.2 Technical Issues with BGP-4.....	12
6. The Routing Information Protocol (RIP).....	13
6.1 Technical Issues with RIP.....	13
7. Bidirectional Forwarding Detection (BFD).....	14
7.1 Technical Issues with BFD.....	14
8. Security Considerations.....	16
9. Acknowledgements.....	16
10. IANA Considerations.....	16
11. References.....	16
11.1 Normative References.....	16
11.2 Informative References.....	17
Contributor's Address.....	19
Author's Addresses.....	19

[1](#). Problem Statement

Protocols, such as OSPF version 2 [[RFC2328](#)], version 3 [[RFC5340](#)], IS-IS [[RFC1195](#)], BGP-4 [[RFC4271](#)] and BFD [[RFC5880](#)], employ various mechanisms to create a cryptographic digest of each transmitted protocol packet. Traditionally, these digests are the results of a one-way hash algorithm, such as Message Digest 5 (MD5) [[RFC1321](#)], across the contents of the packet being transmitted. A secret key is used as the hash base (or seed). The digest is then recomputed by the receiving router, using the same key as the original router used to create the hash, then compared with the transmitted digest to verify:

- o That the router originating this packet is authorized via the shared key mechanism to peer with the local router, and exchange routing data. The implicit trust of routing protocol exchange

protected by a shared secret is intended to protect against the injection of falsely generated routing data into the routing system by unauthorized systems.

- o That the data has not been altered in transit between the two neighboring routers.

Digest verification schemes are not intended to protect the confidentiality of information being exchanged between routers. The information (entries in the routing table) is potentially available through other mechanisms; Moreover, access to the physical media between two routers exchanging routing data, will confer the ability to capture or otherwise discover the contents of the routing tables in those routers.

Authentication mechanisms defined today have notable limitations:

- o Manual configuration of shared secret keys, especially in large Manral, et. al. [Page 3]

networks and between networks, poses a major management problem. In many cases it is challenging to replace keys without significant coordination or disruption.

- o In some cases, when manual keys are configured, some forms of replay protection are no longer possible , allowing the routing protocol to be attacked through the replay of captured routing messages.
- o The MD5 digest algorithm was not designed to be used in the way most routing protocols are using it which has potentially serious future implications.

This document outlines some of the problems with manual keying of these cryptographic algorithms.

1.1 MD5 Pre-Image vs Collision Attacks

A preimage attack (An attempt to find new data with the same hash value) would enable someone to find an input message that causes a hash function to produce a particular output. In contrast, a collision attack finds two messages with the same hash, but the attacker can't pick what the message will be. Feasible collision attacks against MD4, MD5, HAVAL-128, and RIPEMD have been documented in [[CRYPTO-04](#)].

The ability to produce a collision does not currently introduce any obvious or known attacks on routing protocols. Pre-image attacks have the potential to cause problems in the future albeit due to the message length there are serious limitations to the feasibility of mounting such an attack.

Protocols themselves have some built-in protection against collision attacks. This is because a lot of values for fields in a protocol packet are invalid or will produce an unusable packet. For example, in OSPF the Link State Advertisement (LSA) type can be from 1 to 11. Any other value in the field will result in the packet being discarded.

Assume two packets M and M' are generated which have the same hash. The above condition will further reduce the ability to produce a message which is also a correct message from the protocol perspective, as a lot of potential values are themselves not valid.

1.2 Concerns about MD5 and the SHA-1 Algorithm

There are published concerns about the overall strength of the MD5 algorithm ([[Dobb96a](#)], [[Dobb96b](#)], [[Wang04](#)]). While those published concerns apply to the use of MD5 in other modes (e.g., use of MD5 X.509v3/PKIX digital certificates), they are not an attack upon Keyed-MD5, and Hash-based Message Authentication Code MD5 (HMAC-MD5) Manral, et. al. [Page 4]

which is what the current routing protocols have specified. There are also published concerns about the Secure Hash Algorithm (SHA) algorithm ([[Wang05](#)], [[Philip01](#)], [[Prav01](#)], [[Prav02](#)], [[Arjen05](#)]) and also concerns about the MD5 and SHA algorithms in the HMAC [[RFC2104](#)] mode ([[RR07](#)], [[RR08](#)]). National Institute of Standards and Technology (NIST) will be supporting HMAC-SHA-1 even after 2010 [[NIST-HMAC-SHA](#)], whereas it will drop support for SHA-1 as a hash function. NIST also recommends application and protocol designers to move to SHA-2 family of hash functions (i.e., SHA-224, SHA-256, SHA-384 and SHA-512) for all new applications and protocols.

However, as explained above such attacks are currently not applicable to the routing protocols. Separately, some organisations (e.g., the US government) prefer NIST algorithms, such as the SHA family, over other algorithms (like MD5) for local policy reasons.

2. Open Shortest Path First Version 2 (OSPFv2)

OSPF [[RFC2328](#)] describes the use of an MD5 digest with OSPF packets.

MD5 keys are manually configured. The OSPF packet Header includes an authentication type field as well as 64 bits of data for use by the appropriate authentication scheme. OSPF also provides for a non-decreasing sequence number to be included in each OSPF protocol packet to protect against replay attacks.

[2.1](#) Management Issues with OSPFv2

According to the OSPF specification [[RFC2328](#)], digests are applied to packets transmitted between adjacent neighbors, rather than being applied to the routing information originated by a router (digests are not applied at the LSA level, but rather at the packet level). [[RFC2328](#)] states that any set of OSPF routers adjacent across a single link may use a different key to build MD5 digests than the key used to build MD5 digests on any other link. Thus, MD5 keys may be configured, and changed, on a per-link basis in an OSPF network.

OSPF does not specify a mechanism to negotiate keys, nor does it specify any mechanism to negotiate the hash algorithms to be used.

With the proliferation of the number of hash algorithms, as well as the need to continuously upgrade the algorithms, manually configuring the information becomes very tedious. It should also be noted that rekeying OSPF requires coordination among the adjacent routers.

[2.2](#) Technical Issues with OSPFv2

While OSPF provides relatively strong protection through the inclusion of MD5 digests, with additional data and a sequence numbers in transmitted packets, there are still attacks against OSPF:

- o The sequence number is initialized to zero when forming an adjacency with a newly discovered neighbor. When an adjacency is brought down the sequence number is also set to zero. If the cryptographically protected packets of a router that is brought down (for administrative or other reasons) are replayed by a malicious router, traffic could be forced through the malicious router. A malicious router might then induce routing loops, intercept or blackhole the traffic.
- [Page 5]

- o OSPF allows multiple packets with the same sequence number. This means that the possibility exists to replay the same packet many times before the next legitimate packet is sent. An attacker may resend the same packet repeatedly until the next hello packet

is transmitted and received. The Hello interval which is unknown determines the attack window.

- o OSPF does not require the use of any particular hash algorithm, however the use of only MD5 digests for authentication and replay protection is specified in the document. Most OSPF implementations only support MD5 in addition to Null and Simple Password authentication.

Recently, limitations in collision-resistance properties of the MD5 and SHA-1 hash functions have been discovered; [[RFC4270](#)] summarizes the discoveries. There have been attacks against the use of MD5 as a hash; while these attacks do not directly apply to the use of MD5 in routing protocols, it is prudent to have other options available. For this reason the general use of these algorithms should to be discouraged and for this reason [[RFC5709](#)] adds support for using SHA-1 and SHA-2 with the HMAC construct for OSPF.

- o OSPF on a broadcast network shares the same key between all neighbors on that broadcast network. Some OSPF packets are sent to a multicast address.

Spoofing by any malicious neighbor possessing credentials or replayable packets is therefore very easy. Possession of the key itself is used as an identity validation and no other identity check is used. A malicious neighbor could send a packet forging the identity as being from an other neighbor. There would be no way in which the victim could distinguish the identity of the packet sender.

- o In some OSPF implementations neighbors on broadcast, non-broadcast multi-access (NBMA) and point-to-multipoint networks are identified by the IP address in the IP header. The IP header is not covered by the MAC in the cryptographic authentication scheme as described in [RFC 2328](#), and an attack can be made to exploit this omission.

Assume the following scenario.

R1 sends an authenticated HELLO to R2. This HELLO is captured and replayed back to R1. The source IP in the IP header of the replayed packet is changed to that of R2.

R1, not finding itself in HELLO would deduce that the connection is not bidirectional and would bring down the adjacency.

3. Open Shortest Path First Version 3(OSPFv3)

OSPFv3 [[RFC5340](#)] relies on the IP Authentication Header (AH) [[RFC4302](#)] and the IP Encapsulating Security Payload (ESP) [[RFC4303](#)] to cryptographically sign routing information passed between routers.

When using ESP, the null encryption algorithm [[RFC2410](#)] is used, so the data carried in the OSPFv3 packets is signed, but not encrypted. This provides data origin authentication for adjacent routers, and data integrity which gives the assurance that the data transmitted by a router has not changed in transit. However it does not provide confidentiality of the information transmitted.

Authentication/Confidentiality for OSPFv3 [[RFC4552](#)] mandates the use of ESP with null encryption for authentication and also does encourage the use of confidentiality to protect the privacy of the routing information transmitted, using ESP encryption. It however only specifies the use of manual keying of routing information as discussed in the following section.

3.1 Management Issues with OSPFv3

The OSPFv3 security document - Authentication/Confidentiality for OSPFv3 [[RFC4552](#)] discusses, at length, the reasoning behind using manually configured keys, rather than some automated key management protocol such as IKEv2 [[RFC4306](#)]. The primary problem is the lack of suitable key management mechanism, as OSPF adjacencies are formed on a one-to-many basis and most key management mechanisms are designed for a one-to-one communication model. This forces the system administrator to use manually configured security associations (SAs) and cryptographic keys to provide the authentication and, if desired, confidentiality services.

Regarding replay protection [[RFC4552](#)] states that:

As it is not possible as per the current standards to provide complete replay protection while using manual keying, the proposed solution will not provide protection against replay attacks.

The primary administrative issue with manually configured SAs and keys in the OSPFv3 case is the management issues, maintaining shared sets of keys on all routers within a network. As with OSPFv2 rekeying is an infrequent event requiring coordination. [RFC4552] does not require that all OSPFv3 routers have the same key configured for every neighbor, so each set of neighbors connected to a given link could have a different key configured. While this makes it easier to change the keys, by forcing the system administrator to only change the keys on the routers on a single link, the process of manual configuration for all the routers in a network to change the keys used for OSPFv3 digests and confidentiality on a periodic basis can be difficult.

[3.2](#) Technical Issues with OSPFv3

The primary technical concern with the current specifications for OSPFv3 is that when manual security association (SA) and key management is used as [RFC4302] specifies, in [section 3.3.2](#), Sequence Number Generation: "The sender assumes anti-replay is enabled as a default, unless otherwise notified by the receiver (see 3.4.3) or if the SA was configured using manual key management." Replaying OSPFv3 packets can induce several failures in a network, including:

- o Replaying hello packets with an empty neighbor list can cause all the neighbor adjacencies with the sending router to be reset, disrupting network communications.
- o Replaying hello packets from early in the designated router election process on broadcast links can cause all the neighbor adjacencies with the sending router to be reset, disrupting network communications.
- o Replaying database description (DB-Description) packets can cause all FULL neighbor adjacencies with the sending router to be reset, disrupting network communications.
- o Replaying link state request (LS-Request) packets can cause all FULL neighbor adjacencies with the sending router to be reset, disrupting network communications.
- o Capturing a full adjacency process (from two-way all the way to FULL state), and then replaying this process when the router is no longer attached can cause a false adjacency to be formed, allowing an attacker to attract traffic.
- o OSPFv3 on a broadcast network shares the same key between all neighbors on that network. Some OSPF packets are sent to a multicast address.

Spoofing by a malicious neighbor is very easy. Possession of the

identity check used. A neighbor could send a packet specifying the packet came from some other neighbor and there would be no way in which the attacked router could figure out the identity of the packet sender.

[4.](#) Intermediate System to Intermediate System Routing Protocol (IS-IS)

Integrated IS-IS [[RFC1195](#)] uses HMAC-MD5 (Hashed Message Authentication Code MD5) authentication with manual keying, as described in [[RFC5304](#)] and has recently been extended to provide support for using the HMAC construct along with the Secure Hash Algorithm (SHA) [[RFC5310](#)] family of cryptographic hash functions. There is no provision within IS-IS to encrypt the body of a routing protocol message.

[4.1](#) Management Issues with IS-IS

[[RFC5304](#)] states that each Link State Packet (LSP) generated by an intermediate system is signed with the HMAC-MD5 algorithm using a key manually defined by the network administrator. Since authentication is performed on the LSPs transmitted by an intermediate system, rather than on the packets transmitted to a specific neighbor, it is implied that all the intermediate systems within a single flooding domain must be configured with the same key for authentication to work correctly.

The initial configuration of manual keys for authentication within an IS-IS network is simplified by a state where LSPs containing HMAC-MD5/HMAC-SHA authentication TLVs are accepted by intermediate systems without the keys, but the digest is not validated. Once keys are configured on all routers, changing those keys becomes much more difficult.

IS-IS [[RFC1195](#)] does not specify a mechanism to negotiate keys, nor does it specify any mechanism to negotiate the hash algorithms to be used.

With the proliferation of available hash algorithms, as well as the need to upgrade the algorithms, manual configuration requires coordination among intermediate systems which can become tedious.

[4.2](#) Technical Issues with IS-IS

[RFC5304] states: "This mechanism does not prevent replay attacks, however, in most cases, such attacks would trigger existing mechanisms in the IS-IS protocol that would effectively reject old information."

The few cases where existing mechanisms in the IS-IS protocol would not effectively reject old information is the case of Hello packets or the Intermediate System to Intermediate System Hellos (IIHs) that Manral, et. al. [Page 9]

Internet-Draft

Routing Protocol Protection Issues

June 2010

are used to discover neighbors, and the Sequence Number Packets(SNPs).

As described in IS-IS [[RFC1195](#)], a list of known neighbors is included in each Hello transmitted by an intermediate system, to ensure two-way communications with any specific neighbor before exchanging link state databases.

IS-IS does not provide a sequence number. IS-IS packets are vulnerable to replay attacks; any packet can be replayed at any point of time. So long as the keys used are the same, protocol elements that would not be rejected will affect existing sessions.

A hello packet containing a digest within a TLV, and an empty neighbor list, could be replayed, resulting in all adjacencies with the original transmitting intermediate system to be restarted.

A replay of an old Complete Sequence Number Packet (CSNP) could cause LSPs to be flooded, resulting in an LSP storm.

IS-IS specifies the use of the HMAC-MD5 and HMAC-SHA-1 to protect IS-IS packets.

IS-IS does not have a notion of Key ID. During key rollover, each message received has to be checked for integrity against all keys that are valid. A Denial of Service (DoS) attack may be induced by sending IS-IS packets with random hashes. This will cause the IS-IS packet to be checked for authentication with all possible keys, increasing the amount of processing required. This issue however has been fixed in the recent [[RFC5310](#)] which introduces the concept of Key IDs in IS-IS.

Recently, limitations in collision-resistance properties of the MD5 and SHA-1 hash functions have been discovered; [[RFC4270](#)] summarizes the discoveries. There have been attacks against the use of MD5 as a hash; while these attacks do not directly apply to

the use of HMAC-MD5 in IS-IS, it is prudent to have other options available. For this reason the general use of these algorithms should to be discouraged and [[RFC5310](#)] adds support for using HMAC-SHA with IS-IS.

IS-IS on a broadcast network shares the same key between all neighbors on that network.

This makes spoofing by a malicious neighbor easy since IS-IS packets are sent to a link layer multicast address. Possession of the Key itself is used as an authorization check. A neighbor could send a packet spoofing the identity of a neighbor and there would be no way in which the attacked router could discern the identity of the malicious packet sender.

Manral, et. al.

[Page 10]

Internet-Draft

Routing Protocol Protection Issues

June 2010

The Remaining Lifetime field in the LSP is not covered by the authentication. An IS-IS router can receive its own self generated LSP segment with zero lifetime remaining. In that case, if it has a copy with non-zero lifetime, it purges that LSP i.e., it increments the current sequence number and floods all the segments again. This is much worse in IS-IS compared to OSPF, because there is only one LSP other than the pseudonode LSPs for the LANs on which it is the Designated Intermediate System (DIS).

This way an attacker can force the router to flood all segments, potentially a large number if the number of routes is large. It also causes the sequence number of all the LSPs to increase fast. If the sequence number increases to the maximum (0xFFFFFFFF), the IS-IS process must shut down for around 20+ minutes (the product of MaxAge + ZeroAgeLifetime) to allow the old LSPs to age out of all the router databases.

5. Border Gateway Protocol (BGP-4)

BGP-4 [[RFC4271](#)] uses TCP [[RFC0793](#)] for transporting routing information between BGP speakers which have formed an adjacency.

[RFC2385] describes the use of TCP MD5 digest option for providing packet origin authentication and data integrity protection of BGP packets. [[RFC3562](#)] provides suggestions for choosing the key length of the ad-hoc keyed-MD5 mechanism specified in [[RFC2385](#)]. There is no provision for confidentiality for any of these BGP messages.

The IETF has recently obsoleted TCP-MD5 [[RFC2385](#)] with a new TCP

Authentication Option [[TCP-AO](#)]. [[TCP-AO](#)] specifies the use of stronger Message Authentication Codes (MACs), protects against replays even for long lived TCP connections, and provides more details on the association of security with TCP connections than TCP-MD5. It allows rekeying during a TCP connection, assuming that an out-of-band protocol or manual mechanism provides the new keys. Note that TCP MD5 does not preclude rekeying during a connection, but does not require its support either. Further, TCP-AO supports key changes with zero segment loss, whereas key changes in TCP MD5 can lose segments in transit during the changeover or require trying multiple keys on each received segment during key use overlap because it lacks an explicit key ID. Although TCP recovers lost segments through retransmission, loss can have a substantial impact on performance.

This document however still covers only TCP-MD5 as all current deployments are still using BGP with TCP-MD5 and have not upgraded to [[TCP-AO](#)]. There isn't enough operational experience present to evaluate the technical and management issues with this proposal yet.

Compared to previously described IGP protocols, BGP has additional exposure due to the nature of the environment where it is typically used, namely between autonomous networks (under different

Manral, et. al.

[Page 11]

administrative control). While routers running interior gateway protocols may all be configured with the same administrative authority, two BGP peers may be in different administrative domains, having different policies for key strength, rollover frequency, etc. An autonomous system must often support a large number of keys at different BGP boundaries, as each connecting AS represents a different administrative entity. In practice once set, shared secrets between BGP peers are rarely if ever changed.

[5.1](#) Management Issues with BGP-4

Each pair of BGP speakers forming a peering may have a different MD5 shared key facilitating the independent configuration and changing of keys across a large scale network. Manual configuration and maintenance of cryptographic keys across all BGP sessions is a challenge in any large scale environment.

Most BGP implementations will accept BGP packets with a bad digest up to the hold interval negotiated between BGP peers at peering startup, in order to allow for MD5 keys to be changed with minimal impact on operation of the network. This technique does, however, allow some short period of time, during which an attacker may inject BGP packets

with false MD5 digests into the network and can expect those packets to be accepted, even though the MD5 digest is not valid.

[5.2](#) Technical Issues with BGP-4

BGP relies on TCP [[RFC0793](#)] for transporting data between BGP speakers. BGP can rely on TCP's protection against data corruption and replay to preclude replay attacks against BGP sessions. A great deal of research has gone into the feasibility of an attacker overcoming these protections, including [[TCP-WINDOW](#)] and [[Conv01](#)]. Most router and Operating System (OS) vendors have modified their TCP implementations to resolve the security vulnerabilities described in these references, where possible.

As mentioned earlier, MD5 is vulnerable to collision attacks, and can be attacked through several means, such as those explored in [[Wang04](#)].

Though it can be argued that the collision attacks do not have a practical application in this scenario, the use of MD5 should be discouraged.

Routers performing cryptographic processing of packets in software may be exposed to additional opportunities for DoS attacks. An attacker may be able to transmit enough spoofed traffic with false digests that the router's processor and memory resources are consumed, causing the router to be unable to perform normal processing. This exposure is particularly problematic between routers not under unified administrative control.

Manral, et. al.

[Page 12]

[6](#). The Routing Information Protocol (RIP)

The initial version of RIP was specified in STD34 [[RFC1058](#)]. This version did not provide for any authentication or authorization of routing data, and thus was vulnerable to any of a number of attacks against routing protocols. This limitation was one reason why this protocol was moved to Historic status [[RFC1923](#)].

RIPv2, originally specified in [[RFC1388](#)], then [[RFC1723](#)], was finalized in STD56 [[RFC2453](#)]. This version of the protocol provides for authenticating packets with a digest. The details thereof have initially been provided in "RIP-2 MD5 Authentication" [[RFC2082](#)]; "RIPv2 Cryptographic Authentication" [[RFC4822](#)] obsoletes [[RFC2082](#)] and adds details of how the SHA family of hash algorithms can be used

to protect RIPv2. [\[RFC2082\]](#) only specified the use of Keyed MD5.

[6.1](#) Technical Issues with RIP

- o The sequence number used by a router is initialized to zero, at startup, and is also set to zero whenever the neighbor is brought down. If the cryptographically protected packets of a router that is brought down (for administrative or other reasons) are stored by a malicious router, the new router could replay the packets from the previous session thus forcing traffic through the malicious router. Dropping of such packets by the router could result in blackholes. Also forwarding wrong packets could result in routing loops.
- o RIPv2 allows multiple packets with the same sequence number. This could mean the same packet may be replayed many times before the next legitimate packet is sent. An attacker may resend the same packet repeatedly until the next Hello packet is transmitted and received, which means the hello interval therefore determines the attack window.
- o RIPv2 [\[RFC2453\]](#) does not specify the use of any particular hash algorithm. Currently, RIP implementations support keyed MD5 [\[RFC2082\]](#) and [\[RFC4822\]](#) adds support for using HMAC-SHA for RIP.
- o RIPv2 Cryptographic Authentication [\[RFC4822\]](#) does not cover the UDP and the IP headers. It is therefore possible for an attacker to modify some fields in the above headers without routers becoming aware of it.

There is limited exposure to modification of the UDP header as the RIP protocol uses only it to compute the length of the RIP packet. Changes introduced in the UDP header would cause RIP authentication to fail the RIP authentication, limiting exposure.

RIP uses the source IP address from the IP header to determine
Manral, et. al. [Page 13]

which RIP neighbor it has learnt the RIP Update from. Changing the source IP address can be used by an attacker to disrupt the RIP routing sessions between two routers R1 and R2, as shown in the following examples:

Scenario 1:

R1 sends an authenticated RIP message to R2 with a cryptographic

sequence num X.

The attacker then needs a higher sequence number packet originated by R2 either from this session, or from some earlier session.

The attacker can then replay this packet to R2 by changing the source IP to that of R1.

R2 would then no longer accept any more RIP Updates from R1 as those would have a lower cryptographic sequence number. After 180 seconds (or less), R2 would consider R1 timed out and bring down the RIP session.

Scenario 2:

R1 announces a route with cost C1 to R2. This packet can be captured by an attacker. Later, if this cost changes and R1 announces this with a different cost C2, the attacker can replay the captured packet, modifying the source IP to a new arbitrary IP address thereby masquerading as a different router.

R2 will accept this route and the router as a new gateway, and R2 would then use the non existent router as a next hop for that network. This would only be effective if the cost C1 is less than C2.

7. Bidirectional Forwarding Detection (BFD)

BFD is specified in the document [[RFC5880](#)]. Extensions to BFD for Multi-hop [[RFC5883](#)] and single hop [[RFC5881](#)] are defined for IPv4 and IPv6. It is designed to detect failure with the forwarding plane next hop.

BFD base specifies an optional authentication mechanism which can be used by receiver of a packet to be able to authenticate the source of the packet. It relies on the fact that the keys are shared between the peers and no mechanism is defined for the actual Key generation.

7.1 Technical Issues with BFD

- o The level of security provided is based on the Authentication Type used. However the authentication algorithms defined are MD5 or

Manral, et. al.

[Page 14]

be vulnerable to collision attacks.

- o BFD specification mentions mechanisms to allow for the change of authentication state based on the state of a received packet. This can cause a denial of service attack where a malicious authenticated packet (stored from a past session) can be relayed over a session which does not use authentication. This causes one end to assume that authentication is enabled at the other end and hence the BFD adjacency is dropped. This would be a harder attack to put forth when meticulously keyed authentication is in use.
- o BFD works on microsecond timers. When malicious packets are sent at very small duration of time, with the authentication bit set, it can cause a DoS attack.
- o BFD allows a mode called the echo mode. Echo packets are not defined in the BFD specification, though they can keep the BFD session up. There are no guidelines on the properties of the echo packets beyond the choice of the source and destination addresses, and while the BFD specification recommends applying security mechanisms to prevent spoofing of these packets, there are no guidelines on what type of mechanisms are appropriate.

Any security issues in the echo mode will directly affect the BFD protocol and session states and hence the network stability. The potential effects and remedies as understood today are somewhat limited, however. For instance, any replay attacks would be indistinguishable from normal forwarding of the tested router. An attack would still cause a faulty link to be believed to be up, but there is little that can be done about it. However, if the echo packets are guessable it may be possible to spoof from an external source and cause BFD to believe that a one-way link is really bidirectional. As a result it is important that the echo packets contain random material that is also checked upon reception.

- o BFD packets can be sent at millisecond intervals (the protocol uses timers at microsecond intervals). When using authentication this can cause frequent Sequence Number wrap-around as a 32-bit sequence number is used, thus considerably reducing the security of the authentication algorithms.
- o Recently, limitations in collision-resistance properties of the MD5 and SHA-1 hash functions have been discovered; [\[RFC4270\]](#) summarizes the discoveries. There have been attacks against the use of MD5 as a hash; while these attacks do not directly apply to the use of HMAC-MD5 and keyed SHA-1 in BFD, it is prudent to have other options available. Such attacks do not mean that BFD using SHA-1 for authentication is at risk. However, it does mean that SHA-1 should be replaced as soon as possible and should not be used

for new applications.

It should be noted that if SHA-1 is used in the Hashed Message Authentication Code (HMAC) [[RFC2104](#)] construction then collision attacks currently known against SHA-1 do not apply. The new attacks on SHA-1 have no impact on the security of HMAC-SHA-1.

There are already proposals [[GEN-BFD-AUTH](#)] that add support for HMAC with SHA-1 and SHA-2 family of hash functions for BFD.

[8](#). Security Considerations

This draft outlines security issues arising from the current methodology for manual keying of various routing protocols. No specific changes to routing protocols are proposed in this draft, likewise no new security requirements result.

[9](#). Acknowledgements

We would like to acknowledge Sam Hartman, Ran Atkinson, Stephen Kent and Brian Weis for their initial comments on this draft. Thanks to Merike Kaeo and Alfred Hoenes for reviewing many sections of the draft and providing lot of useful comments.

[10](#). IANA Considerations

This document places no requests to IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

[11](#). References

[11.1](#) Normative References

- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), September 1981.
- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", [RFC 1195](#), December 1990.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), April 1998.

[RFC2385] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", [RFC 2385](#), August 1998.

[RFC2453] Malkin, G., "RIP Version 2", [RFC 2453](#), November 1998.

Manral, et. al.

[Page 16]

Internet-Draft

Routing Protocol Protection Issues

June 2010

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and Lindem, A., "OSPF for IPv6", [RFC 5340](#), July 2008.

[RFC5304] Li, T. and Atkinson, R., "Intermediate System to Intermediate System (IS-IS) Cryptographic Authentication", [RFC 5304](#), October 2008.

[RFC5310] Bhatia, M., et. al., "IS-IS Generic Cryptographic Authentication", [RFC 5310](#), February 2009

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.

[RFC4302] Kent, S., "IP Authentication Header", [RFC 4302](#), December 2005.

[RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", [RFC 4303](#), December 2005.

[RFC4552] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", [RFC 4552](#), January 2006.

[RFC4822] Atkinson, R. and Fanto, M., "RIPv2 Cryptographic Authentication", [RFC 4822](#), February 2007.

[11.2](#) Informative References

[Arjen05] Arjen K. Lenstra, "Further progress in Hashing cryptanalysis", Lucent Bell Laboratories, February 26, 2005.

[Conv01] Convery, S. and M. Franz, "BGP Vulnerability Testing: Separating Fact from FUD v1.00", June 2003.

[CRYPTO-04] Xiaoyun Wang, Xuejia Lai, Dengguo Feng and Hongbo Yu, "Collisions for hash functions MD4, MD5, HAVAL-128, and RIPEMD", Crypto 2004 Rump Session.

[Dobb96a] Dobbertin, H., "Cryptanalysis of MD5 Compress", Technical

Report, 2 May 1996. (Presented at the Rump Session of EuroCrypt 1996.)

[Dobb96b] Dobbertin, H, "The Status of MD5 After a Recent Attack", CryptoBytes, Vol. 2, No. 2, Summer 1996.

[GEN-BFD-AUTH] Bhatia, Manav and Manral, Vishwas, "BFD Generic Cryptographic Authentication", Work in Progress

[NIST-HMAC-SHA] "NISTs Policy on Hash Functions", 2006. Available
Manral, et. al. [Page 17]

Internet-Draft

Routing Protocol Protection Issues

June 2010

online at <http://csrc.nist.gov/groups/ST/hash/policy.html>

[Philip01] Philip Hawkes¹, Michael Paddon¹, and Gregory G. Rose, "On Corrective Patterns for the SHA-2 Family", Qualcomm Australia

[Prav01] Praveen Gauravaram, et. al., "Collision Attacks on MD5 and SHA-1: Is this the "Sword of Damocles" for Electronic Commerce?", Information Security Institute (ISI), Queensland University of Technology (QUT), Australia

[Prav02] Praveen Gauravaram, et. al., "Some thoughts on Collision Attacks in the Hash Functions Md5, SHA-0 and SHA-1", Information Security Institute (ISI), Queensland University of Technology (QUT), Australia

[RFC1058] Hedrick, C., "Routing Information Protocol", [RFC 1058](#), June 1988.

[RFC1321] Rivest, R., "The MD5 Message-Digest Algorithm", [RFC 1321](#), April 1992.

[RFC1388] Malkin, G., "RIP Version 2 Carrying Additional Information", [RFC 1388](#), January 1993.

[RFC1723] Malkin, G., "RIP Version 2 - Carrying Additional Information", STD 56, [RFC 1723](#), November 1994.

[RFC1923] Halpern, J. and Bradner, S., "RIPv1 Applicability Statement for Historic Status", [RFC 1923](#), March 1996.

[RFC2082] Baker, F. and Atkinson, R., "RIP-2 MD5 Authentication", [RFC 2082](#), January 1997

- [RFC2104] Krawczyk, H., et. al., "HMAC: Keyed-Hashing for Message Authentication", [RFC 2104](#), February 1997
- [RFC2410] Kent, S. and Glenn, R., "The NULL Encryption Algorithm and Its Use With IPsec", [RFC 2410](#), November 1998
- [RFC3562] Leech, M., "Key Management Considerations for the TCP MD5 Signature Option", [RFC 3562](#), July 2003.
- [RFC4270] Hoffman, P. and B. Schneier, "Attacks on Cryptographic Hashes in Internet Protocols", [RFC 4270](#), November 2005.
- [RFC4306] Kaufman, C., "The Internet Key Exchange (IKEv2) Protocol", [RFC 4306](#), December 2005.
- [RFC5709] Bhatia, M. et. al, "OSPFv2 HMAC-SHA Cryptographic Manral, et. al. [Page 18]
-

Internet-Draft Routing Protocol Protection Issues June 2010

- Authentication", [RFC 5709](#), October 2009.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection", [RFC 5880](#), June 2010
- [RFC5881] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", [RFC 5881](#), June 2010
- [RFC5883] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for Multihop Paths", [RFC 5883](#), June 2010
- [RR07] Rechberger, C. and V. Rijmen, "On Authentication with HMAC and Non-random Properties", Financial Cryptography and Data Security, Lecture Notes in Computer Science, Volume 4886/2008, Springer-Verlag, Berlin, December 2007.
- [RR08] Rechberger, C. and V. Rijmen, "New Results on NMAC/HMAC when Instantiated with Popular Hash Functions", Journal of Universal Computer Science, Volume 14, Number 3, pp. 347-376, 1 February 2008.
- [TCP-WINDOW] Watson, T., "TCP Reset Spoofing", October 2003.
- [TCP-A0] Touch, J., Mankin, A. and Bonica, R., "The TCP Authenticaion Option", Work in Progress

- [Wang04] Wang, X., et alia, "Collisions for Hash Functions MD4, MD5, HAVAL-128, and RIPEMD", August 2004, IACR, <http://eprint.iacr.org/2004/199>
- [Wang05] Wang, X., et alia, "Finding Collisions in the Full SHA-1" Proceedings of Crypto 2005, Lecture Notes in Computer Science, Volume 3621, pp. 17-36, Springer-Verlag, Berlin, August 31, 2005.

Contributor's Address

Sue Hares
NextHop
USA
Email: shares@nexthop.com

Author's Addresses

Vishwas Manral
IP Infusion
Manral, et. al.

[Page 19]

Internet-Draft

Routing Protocol Protection Issues

June 2010

Almora, Uttarakhand
India
Email: vishwas@ipinfusion.com

Manav Bhatia
Alcatel-Lucent
Bangalore
India
Email: manav.bhatia@alcatel-lucent.com

Joel P. Jaeggli
Check Point Software
Email: jjaeggli@checkpoint.com

Russ White
Cisco Systems
RTP North Carolina
USA
Email: riw@cisco.com

