Network Working Group **Vishwas Manral**
Internet-Draft **IP Infusion**
Expires: August 2008 **Russ White**
Intended Status: Informational Cisco Systems
Manav Bhatia
Alcatel-Lucent

Issues with existing Cryptographic Protection Methods for Routing Protocols

draft-manral-rpsec-existing-crypto-05.txt

Status of this Memo

Abstract

Routing protocols are designed to use cryptographic mechanisms to authenticate data being received from a neighboring router to ensure that it has not been modified in transit, and actually originated from the neighboring router purporting to have originating the data. Most of the cryptographic mechanisms defined to date rely on hash algorithms applied to the data in the routing protocol packet, which means the data is transported, in the clear, along with a signature based on the data itself. These mechanisms rely on the manual configuration of the keys used to seed, or build, these hash based signatures. This document outlines some of the problems with manual

keying of these cryptographic algorithms.

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

Routing protocols, such as OSPF [RFC2740] [RFC2328], IS-IS [RFC1195], and BGP-4 [RFC4271], rely on various mechanisms to create a cryptographic digest of each transmitted routing protocol. Traditionally, these digests are the results of a hash algorithm, such as MD5 [RFC1321], across the contents of the packet being transmitted, using a secret key as the hash base (or seed). These digests are recomputed by the receiving router, using the same key as the originating router used to create the hash, and compared with the transmitted digest to verify:

- That the router originating this piece of data is authorized to peer with the local router, and to transmit routing data. This

generally protects against falsely generated routing data being injected into a routing system by rogue systems.

- That the data has not been changed while transiting between the two neighboring routers.

These sorts of authentication methods are not generally used to protect the confidentiality of information being exchanged between routers, since this information (entries in the routing table) is generally freely available in many other context; if anyone has access to the physical media between two routers exchanging routing data, they will also probably have other ways to capture or otherwise discover the contents of the routing tables in those routers.

The main problems with the authentication mechanisms defined today revolve around:

- Manual configuration of shared secret keys, especially in large scale networks, poses a major management problem, especially as there is generally no way to gracefully move from one secret key to another.

- In some cases, when manual keys are configured, some forms of replay protection are disabled, allowing the routing protocol to be attacked.

In fact, the MD5 digest algorithm was not designed to be used in the way most routing protocols are using it, which can lead to serious security implications in the future.

A preimage attack 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 hash will be. Feasible collision attacks against MD4, MD5, HAVAL-128, and RIPEMD were found by the Chinese researcher Xiaoyun Wang with co-authors Dengguo Feng, Xuejia Lai, and Hongbo Yu.

The collision vulnerability does not introduce any obvious or known attacks on routing protocols. However pre-image attacks could cause problems.

Protocols themselves have some built-in protection against collision attacks. A lot of values for a lot of fields in the protocol are invalid. For example, for OSPF the LSA type can be from 1 to 11. Any other value in the field will result in the packet being dropped.

Assume two packets M and M' are generated which have the same hash. The above condition will further reduce the probability of the two messages also being correct messages from the protocol perspective, as a lot of values are themselves not valid.

2. Open Shortest Path First (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 OSPF

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 mechanisms 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.

2.2 Technical Issues with OSPF

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

- The sequence number is initialized to zero when forming an adjacency with a newly discovered neighbor, 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.

- OSPF allows multiple packets with the same sequence number.

This could mean the same packet can 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 determines the attack window.

- OSPF does not specify the use of any particular hash algorithm, however the use of only MD5 is specified in the document. Most OSPF implementations only support MD5.

  Recently, attacks on the collision-resistance property of the MD5 and SHA-1 hash functions have been discovered; [RFC4270] summarizes the discoveries. The attacks on MD5 are practical on any modern computer. For this reason the use of these algorithms needs to be discouraged.

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

  This allows spoofing by any malicious neighbor very easy. Possession of the key itself is used as an identity check. There is no other 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.

- OSPF neighbors on broadcast, NBMA and point-to-multipoint networks are identified by the IP address in the IP header. Because the IP header is not covered by the MAC in the cryptographic authentication scheme as described in RFC 2328, an attack can be made exploiting this vulnerability.

  Assume the following scenario.

  R1 sends an authenticated HELLO to R2. This HELLO is captured and replayed back to R1, changing the source IP in the IP header 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 (OSPFv3)

OSPFv3 [RFC2740] relies on the IP Authentication Header described in [RFC4302] and the IP Encapsulating Payload described in [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 data transmitted by a router has not changed in transit.

However it does not provide confidentiality of the information transmitted. [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.

[RFC4552] describes OSPFv3's use of AH and ESP, and specifies that only manual keying of routing information may be used.

3.1 Management Issues with OSPFv3

The OSPFv3 security document [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 that all current key management mechanisms are designed for a one-to-one correlation of keys, while OSPF adjacencies are formed on a one-to-many basis. This forces the system administrator to use manually configured SAs and cryptographic keys to provide the authentication and, if desired, confidentiality services.

[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 SA's and keys in the OSPFv3 case is the simple management issue of maintaining matching sets of keys on all routers within a network. [RFC4552] does not require that all OSPFv3 routers have the same key configured for every neighbor, so each set of neighbors connected to a single 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 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." Replayed OSPFv3 packets can cause

several failures in a network, including:

- Replaying hello packets with an empty neighbor list can cause all the neighbor adjacencies with the sending router to be reset,

disrupting network communications.

- 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.

- Replaying database description (DB-Description) packets can cause all FULL neighbor adjacencies with the sending router to be reset, disrupting network communications.

- Replaying link state request (LS-Request) packets can cause all FULL neighbor adjacencies with the sending router to be reset, disrupting network communications.

- 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 and black hole traffic.

- OSPFv3 on a broadcast network shares the same key between all neighbors on that network. Some OSPF packets are sent to a multicast address.

  This allows spoofing by any malicious neighbor very easy. Possession of the key itself is used as an identity check. There is no other 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 authentication with manual keying, as described in [RFC3567]. There is no provision within IS-IS to encrypt the body of a routing protocol message.

4.1 Management Issues with IS-IS

[RFC3567] states that each 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 authentication TLVs are accepted, but the digest is not validated. Once an initial set of keys is configured on all routers, however, 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 the number of hash algorithms, as well as

the need to continuously upgrade the algorithms, manually configuring the information becomes very tedious.

4.2
Technical Issues with IS-IS

[RFC3567] 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 (IIHs) used to discover neighbors, and SNP packets.

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. Hence IS-IS packets are liable to replay attacks; any packet can be replayed at any point of time, as long as the keys used are the same.

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

A replay of an old CSNP packets could cause LSPs to be flooded, thus causing an LSP storm.

IS-IS specifies the use of the hash algorithm HMAC-MD5 to protect IS-IS PDUs.

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 DoS attack may be caused by sending IS-IS packets with random hashes. This will cause the IS-IS packet to be checked for authentication with all possible keys, thus increasing the amount of processing required.

Recently, attacks on the collision-resistance property of the MD5 and SHA-1 hash functions have been discovered; [RFC4270] summarizes the discoveries. The attacks on MD5 are practical on any modern computer.

For this reason, the use of these algorithms needs to be discouraged.

HMACs are not susceptible to any known collision-reduction attack.

However, IS-IS should provide a way to upgrade to other stronger algorithms.

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

This makes spoofing by any malicious neighbor very easy since IS-IS PDUs are sent to a link layer multicast address. Possession of the key itself is used as an identity check and no other identity check is performed. 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.

As the lifetime is not covered in 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, as there exists only one LSP other than the pseudonode LSPs for the LANs on which it is the Designated Intermediate System (DIS).

This way an attack can force the router to flood all segments, which can be quite a lot 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 (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 signature option for providing data origin authentication and data integrity protection of these BGP packets, and [RFC3562] gives suggestions for choosing the key length for the ad-hoc keyed-MD5 mechanism specified in [RFC2385]. There is no provision for confidentiality for any of these BGP messages.

This problem is made worse by the nature of the environment where BGP is typically used, between autonomous networks (under different administrative control). While routers running interior gateway protocols may all be configured using the same keys, and have their key rollover policies coordinated or set by the same administrative authority, two BGP peering BGP speakers may be in different administrative domains, with different policies for key strength, rollover times, etc. An autonomous system must often support a large number of keys on different BGP borders, since each connecting AS

represents a different administrative entity.

5.1 Management Issues with BGP-4

Each pair of BGP speakers forming an adjacency may have a different MD5 shared key, facilitating the configuration and changing of keys across a large scale network. Manual configuration and maintenance of cryptographic keys on all routers is a challenge in any large scale environment, however. Most BGP implementations will accept BGP packets with a bad digest for the hold interval negotiated between BGP peers at peering startup, allowing MD5 keys to be changed without impacting the 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 their MD5 digests are not valid.

5.2 Technical Issues with BGP-4

Since BGP relies on TCP [RFC0793] for transporting data between BGP speakers, BGP can rely on TCP's protections against data corruption and replay to prevent replay attacks against BGP sessions. A great deal of research has gone into the difficulty or ease with which an attacker can overcome these protections, including [TCP-WINDOW] and [BGP-ATTACK]. Most implementations of BGP have modified their TCP implementations to resolve the security vulnerabilities described in these references, where possible.

However, as mentioned earlier, MD5 is vulnerable to collision attacks, and can be attacked through several means, such as those explored in [MD5-ATTACK].

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

Routers performing cryptographic processing of packets in software may be easier to attack. An attacker may be able to transmit enough traffic with false digests to a router that the router's processor and memory resources are consumed, causing the router to be unable to perform normal processing. This is particularly problematic at connections to devices not under local administrative control.

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 the various attacks

against routing protocols. This was one of the reasons why this protocol has been moved to Historic status long ago [RFC1923].

RIPv2, originally specified in [RFC1388], then [RFC1723], has been finalized in STD56 [RFC2453]. This version of the protocol provides for authenticating packets by carrying 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 protext RIPv2, whereas [RFC2082] only specified the use of Keyed MD5.

- The sequence number is initialized to zero, at the beginning of time, 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.

- RIPv2 allows multiple packets with the same sequence number. This could mean the same packet can 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 determines the attack window.

- RIPv2 does not specify the use of any particular hash algorithm. Currently, RIP implementations only support keyed MD5 [RFC2082]. MD5 is vulnerable to attacks [MD5-ATTACK].

- RIPv2 Cryptographic Authentication [RFC4822] does not cover the UDP and the IP headers. It is thus possible for an attacker to modify the fields in the above headers without any of the routers getting to know about it.

  There isn't much that can be done by modifying the UDP header as RIP only uses it to compute the length of the RIP packet. Any changes introduced in the UDP header would fail the RIP authentication, and this attack will thus, not work.

  However, RIP uses the source IP address from the IP header to determine the RIP neighbor from which it has learnt the RIP Updates. This 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 merely needs get hold of a higher sequence number packet from the LAN. It could also be a 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 now no longer accept any more RIP Updates from R1 as

those would have a lower cryptographic sequence number. After 180 secs (or less), R2 would time out R1 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 some other cost C2, the attacker can replay the captured packet by modifying the source IP to some new arbitrary IP address. It can this way masquerade as some other router.

R2 will accept this route and the router as a new gateway, and would use the non existent router as a next hop for that network. This would obviously only work if C1 < C2.

7. Security Considerations

This draft outlines security issues arising from the manual keying of cryptographic keys for various routing protocols. No changes to any protocols are proposed in this draft, and no new security requirements result.

8. Acknowledgements

We would like to acknowledge Sam Hartman, Ran Atkinson, Steve 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.

9. IANA Considerations

This document places no requests to IANA.

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

10. References

10.1 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

[RFC2740] Coltun, R., Ferguson, D., and J. Moy, "OSPF for IPv6", RFC 2740, December 1999.

[RFC3567] Li, T. and R. Atkinson, "Intermediate System to Intermediate System (IS-IS) Cryptographic Authentication", RFC 3567, July 2003.

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

10.2 Informative References

[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

[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

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

[TCP-WINDOW] Watson, T., "TCP Reset Spoofing", October 2003.

[MD5-ATTACK] Wang, X. et al., "Collisions for Hash Functions MD4, MD5, HAVAL-128 and RIPEMD", August 2004, http://eprint.iacr.org/2004/199

11. Contributor's Address

Sue Hares
NextHop
USA
Email: shares@nexthop.com

12. Author's Addresses

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

Vishwas Manral
IP Infusion
Almora, Uttarakhand
India

Email: vishwas@ipinfusion.com

Russ White
Cisco Systems

RTP North Carolina
USA
Email: riw@cisco.com

Full Copyright Statement

Intellectual Property