

| | | |
|-----------------------------------|---------------------|--|
| KARP | G. Lebovitz | |
| Internet-Draft | Juniper | |
| Intended status: Informational | January 14, 2010 | |
| Expires: July 18, 2010 | | |

[TOC](#)

Roadmap for Cryptographic Authentication of Routing Protocol Packets on the Wire

draft-lebovitz-karp-roadmap-00

Abstract

In the March of 2006 the IAB held a workshop on the topic of "Unwanted Internet Traffic". The report from that workshop is documented in [RFC 4948 \(Andersson, L., Davies, E., and L. Zhang, "Report from the IAB workshop on Unwanted Traffic March 9-10, 2006," August 2007.\)](#)

[RFC4948]. Section 8.2 of RFC 4948 calls for "[t]ightening the security of the core routing infrastructure." Four main steps were identified for improving the security of the routing infrastructure. One of those steps was "securing the routing protocols' packets on the wire." One mechanism for securing routing protocol packets on the wire is the use of per-packet cryptographic message authentication, providing both peer authentication and message integrity. Many different routing protocols exist and they employ a range of different transport subsystems. Therefore there must necessarily be various methods defined for applying cryptographic authentication to these varying protocols. Many routing protocols already have some method for accomplishing cryptographic message authentication. However, in many cases the existing methods are dated, vulnerable to attack, and/or employ cryptographic algorithms that have been deprecated. This document creates a roadmap of protocol specification work for the use of modern cryptographic mechanisms and algorithms for message authentication in routing protocols. It also defines the framework for a key management protocol that may be used to create and manage session keys for message authentication and integrity. This roadmap reflects the input of both the security area and routing area in order to form a jointly agreed upon and prioritized work list for the effort. This version is actually the fourth version, but is recently renamed from "-kmart-roadmap" to "-karp-roadmap" to follow the new working group name.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

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/ietf/1id-abstracts.txt>.

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

This Internet-Draft will expire on July 18, 2010.

Copyright Notice

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

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 BSD License.

Table of Contents

- [1.](#) Introduction
 - [1.1.](#) Terminology
 - [1.2.](#) Requirements Language
 - [1.3.](#) Scope
 - [1.4.](#) Goals
 - [1.5.](#) Non-Goals
 - [1.6.](#) Audience
- [2.](#) Threats
 - [2.1.](#) Threats In Scope
 - [2.2.](#) Threats Out of Scope
- [3.](#) Categorizing Routing Protocols
 - [3.1.](#) Category: Messaging Transaction Type
 - [3.2.](#) Category: Peer vs. Group Keying
 - [3.3.](#) Category: Update vs. Discovery Protocol
 - [3.4.](#) Security Characterization Vectors
 - [3.4.1.](#) Internal vs. External Operation
 - [3.4.2.](#) Unique versus Shared Keys
 - [3.4.3.](#) Out-of-Band vs. In-line Key Management
- [4.](#) The Roadmap

| | |
|-----------------------|---|
| 4.1. | Work Phases on any Particular Protocol |
| 4.2. | Requirements for Phase 1 Routing Protocols' Security Update |
| 4.3. | Common Framework |
| 4.4. | Work Items Per Routing Protocol |
| 4.5. | Protocols in Categories |
| 4.6. | Priorities |
| 5. | Security Considerations |
| 6. | IANA Considerations |
| 7. | Acknowledgements |
| 8. | Change History (RFC Editor: Delete Before Publishing) |
| 9. | Needs Work in Next Draft (RFC Editor: Delete Before Publishing) |
| 10. | References |
| 10.1. | Normative References |
| 10.2. | Informative References |
| § | Authors' Addresses |

1. Introduction

[TOC](#)

In March 2006 the Internet Architecture Board (IAB) held a workshop on the topic of "Unwanted Internet Traffic". The report from that workshop is documented in [RFC 4948 \(Andersson, L., Davies, E., and L. Zhang, "Report from the IAB workshop on Unwanted Traffic March 9-10, 2006," August 2007.\)](#) [RFC4948]. Section 8.1 of that document states that "A simple risk analysis would suggest that an ideal attack target of minimal cost but maximal disruption is the core routing infrastructure." Section 8.2 calls for "[t]ightening the security of the core routing infrastructure." Four main steps were identified for that tightening:

- *More secure mechanisms and practices for operating routers. This work is being addressed in the OPSEC Working Group.
- *Cleaning up the Internet Routing Registry repository [IRR], and securing both the database and the access, so that it can be used for routing verifications. This work should be addressed through liaisons with those running the IRR's globally.
- *Specifications for cryptographic validation of routing message content. This work will likely be addressed in the SIDR Working Group.
- *Securing the routing protocols' packets on the wire

This document addresses the last bullet, securing the packets on the wire of the routing protocol exchanges. The document addresses Keying and Authentication for Routing Protocols, aka "KARP".

It is unlikely that this document, in its current form, will become an RFC. More likely is that this document will be split up into several smaller documents which may look something like:

- *Scope & Goals sections will likely become part of the KARP WG charter

- *Threat document

- *Requirements document (may be combined with Threat document)

- *Framework document

- *RoutingProtocol Design Team Work Plan document. This would include sections like Work Phases, Priorities, Security Considerations, etc.

For now, the document serves as the catch all for the set of thoughts around the KARP effort. As a working group is formed, decisions will be made about the creation of specific documents.

Editor's Note on "KMART" vs "KARP": The first few versions of this document were called "draft-lebovitz-kmart-roadmap-xx". This went up to -03. Upon the creation of the BoF for IETF76, the IESG requested the name of the effort change so as to avoid any potential trademark issues. The new name of the effort is KARP. Version -03 went out titled "draft-lebovitz-kmart-roadmap-03", so as to avoid last minute confusion at that IETF meeting. This version now changes the "kmart" to "karp" in the title, changes the version counter back to -00, and contains no other changes, and is published as "draft-lebovitz-karp-roadmap-00".

1.1. Terminology

[TOC](#)

Within the scope of this document, the following words, when beginning with a capital letter, or spelled in all capitals, hold the meanings described to the right of each term. If the same word is used uncapitalized, then it is intended to have its common english definition.

PSK Pre-Shared Key. A key used by both peers in a secure configuration. Usually exchanged out-of-band prior to a first connection.

Routing Protocol When used with capital "R" and "P" in this document the term refers the Routing Protocol for which work is being done to provide or enhance its peer authentication mechanisms.

PRF

Pseudorandom number function, or sometimes called pseudorandom number generator (PRNG). An algorithm for generating a sequence of numbers that approximates the properties of random numbers. The sequence is not truly random, in that it is completely determined by a relatively small set of initial values that are passed into the function. An example is SHA-256.

KDF Key derivation function. A particular specified use of a PRF that takes a PSK, combines it with other inputs to the PRF, and produces a result that is suitable for use as a Traffic Key.

Identifier The type and value used by one peer of an authenticated message exchange to signify to the other peer who they are. The Identifier is used by the receiver as a lookup index into a table containing further information about the peer that is required to continue processing the message, for example a Security Association (SA) or keys.

Identity Proof A cryptographic proof for an asserted identity, that the peer really is who they assert themselves to be. Proof of identity can be arranged between the peers in a few ways, for example PSK, raw asymmetric keys, or a more user-friendly representation of asymmetric keys, like a certificate.

Security Association or SA The parameters and keys that together form the required information for processing secure sessions between peers. Examples of items that may exist in an SA include: Identifier, PSK, Traffic Key, cryptographic algorithms, key lifetimes.

KMP Key Management Protocol. A protocol used between peers to exchange SA parameters and Traffic Keys. Examples of KMPs include IKE, TLS, and SSH.

KMP Function Any actual KMP used in the general KARP solution framework

Peer Key Keys that are used between peers as the identity proof. These keys may or may not be connection specific, depending on who they were established, and what form of identity and identity proof is being used in the system.

Traffic Key The actual key used on each packet of a message.

Definitions of items specific to the general KARP framework are described in more detail in the Framework section [Section 4.3 \(Common Framework\)](#).

1.2. Requirements Language

[TOC](#)

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 [RFC2119 \(Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.\)](#) [RFC2119].

When used in lower case, these words convey their typical use in common language, and are not to be interpreted as described in [RFC2119 \(Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.\)](#) [RFC2119].

1.3. Scope

[TOC](#)

Four basic tactics may be employed in order to secure any piece of data as it is transmitted over the wire: privacy (or encryption), authentication, message integrity, and non-repudiation. The focus for this effort, and the scope for this roadmap document, will be message authentication and packet integrity only. This work explicitly excludes, at this point in time, the other two tactics: privacy and non-repudiation. Since the objective of most routing protocols is to broadly advertise the routing topology, routing messages are commonly sent in the clear; confidentiality is not normally required for routing protocols. However, ensuring that routing peers truly are the trusted peers expected, and that no rogue peers or messages can compromise the stability of the routing environment is critical, and thus our focus. The other two explicitly excluded tactics, privacy and non-repudiation, may be addressed in future work.

It is possible for routing protocol packets to be transmitted employing all four security tactics mentioned above using existing standards. For example, one could run unicast, layer 3 or above routing protocol packets through [IPsec ESP \(Kent, S., "IP Encapsulating Security Payload \(ESP\)," December 2005.\)](#) [RFC4303]. This would provide the added benefit of privacy, and non-repudiation. However, router platforms and systems have been fine tuned over the years for the specific processing necessary for routing protocols' non-encapsulated formats. Operators are, therefore, quite reluctant to explore new packet encapsulations for these tried and true protocols.

In addition, at least in the case of BGP and LDP, these protocols already have existing mechanisms for cryptographically authenticating and integrity checking the packets on the wire. Products with these mechanisms have already been produced, code has already been written and both have been optimized for the existing mechanisms. Rather than turn away from these mechanisms, we want to enhance them, updating them to modern and secure levels.

There are two main work phases for this roadmap, and for any Routing Protocol work undertaken as part of this roadmap (discussed further in the [Work Phases \(Work Phases on any Particular Protocol\)](#) section). The first is to enhance the Routing Protocol's current authentication mechanism, ensuring it employs modern cryptographic algorithms and methods for its basic operational model, fulfilling the requirements defined in the [Requirements \(Requirements for Phase 1 Routing Protocols' Security Update\)](#) section, and protecting against as many of the threats as possible as defined in the [Threats \(Threats In Scope\)](#) section. Many of the Routing Protocols' current mechanisms use manual keys, so the first phase updates will focus on shoring up the manual key mechanisms that exist.

The second work phase is to define the use of a key management protocol (KMP) for creating and managing session keys used in the Routing Protocols' message authentication and data integrity functions. It is intended that a general KMP framework -- or a small number of frameworks -- can be defined and leveraged for many Routing Protocols. Therefore, the scope of this roadmap of work includes:

- o Making use of existing routing protocol security protocols, where they exist, and enhancing or updating them as necessary for modern cryptographic best practices,
- o Developing a framework for using automatic key management in order to ease deployment, lower cost of operation, and allow for rapid responses to security breaches, and
- o Specifying the automated key management protocol that may be combined with the bits-on-the-wire mechanisms.

The work also serves as an agreement between the Routing Area and the Security Area about the priorities and work plan for incrementally delivering the above work. This point is important. There will be times when the best-security-possible will give way to vastly-improved-over-current-security-but-admittedly-not-yet-best-security-possible, in order that incremental progress toward a more secure Internet may be achieved. As such, this document will call out places where agreement has been reached on such trade offs.

This document does not contain protocol specifications. Instead, it defines the areas where protocol specification work is needed and sets a direction, a set of requirements, and a relative priority for addressing that specification work.

There are a set of threats to routing protocols that are considered in-scope for this document/roadmap, and a set considered out-of-scope. These are described in detail in the [Threats \(Threats\)](#) section below.

1.4. Goals

The goals and general guidance for this work roadmap follow:

1. Provide authentication and integrity protection for packets on the wire of existing routing protocols
2. Deliver a path to incrementally improve security of the routing infrastructure. The principle of crawl, walk, run will be in place. Routing protocol authentication mechanisms may not go immediately from their current state to a state containing the best possible, most modern security practices. Incremental steps will need to be taken for a few very practical reasons. First, there are a considerable number of deployed routing devices in operating networks that will not be able to run the most modern cryptographic mechanisms without significant and unacceptable performance penalties. The roadmap for any one routing protocol MUST allow for incremental improvements on existing operational devices. Second, current routing protocol performance on deployed devices has been achieved over the last 20 years through extensive tuning of software and hardware elements, and is a constant focus for improvement by vendors and operators alike. The introduction of new security mechanisms affects this performance balance. The performance impact of any incremental step of security improvement will need to be weighed by the community, and introduced in such a way that allows the vendor and operator community a path to adoption that upholds reasonable performance metrics. Therefore, certain specification elements may be introduced carrying the "SHOULD" guidance, with the intention that the same mechanism will carry a "MUST" in the next release of the specification. This gives the vendors and implementors the guidance they need to tune their software and hardware appropriately over time. Last, some security mechanisms require the build out of other operational support systems, and this will take time. An example where these three reasons are at play in an incremental improvement roadmap is seen in the improvement of [BGP's \(Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 \(BGP-4\)," January 2006.\)](#) [RFC4271] security via the update of the TCP Authentication Option ([TCP-AO \(Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option," March 2010.\)](#) [I-D.ietf-tcpm-tcp-auth-opt] effort. It would be ideal, and reflect best common security practice, to have a fully specified key management protocol for negotiating TCP-AO's authentication material, using certificates for peer authentication in the keying. However, in the spirit of incremental deployment, we will first address issues like cryptographic algorithm agility, replay attacks, TCP session resetting in the base TCP-AO protocol before we layer key management on top of it.

3.

The deploy-ability of the improved security solutions on currently running routing infrastructure equipment. This begs the consideration of the current state of processing power available on routers in the network today.

4. Operational deploy-ability - A solutions acceptability will also be measured by how deployable the solution is by common operator teams using common deployment processes and infrastructures. I.e. We will try to make these solutions fit as well as possible into current operational practices or router deployment. This will be heavily influenced by operator input, to ensure that what we specify can -- and, more importantly, will -- be deployed once specified and implemented by vendors. Deployment of incrementally more secure routing infrastructure in the Internet is the final measure of success. Measurably, we would like to see an increase in the number of surveyed respondents who report deploying the updated authentication mechanisms anywhere across their network, as well as a sharp rise in usage for the total percentage of their network's routers.

Interviews with operators show several points about routing security. First, over 70% of operators have deployed transport connection protection via TCP-MD5 on their EBGP [[ISR2008](#)] ([McPherson, D. and C. Labovitz, "Worldwide Infrastructure Security Report," October 2008.](#)) . Over 55% also deploy MD5 on their IBGP connections, and 50% deploy MD5 on some other IGP. The survey states that "a considerable increase was observed over previous editions of the survey for use of TCP MD5 with external peers (eBGP), internal peers (iBGP) and MD5 extensions for IGPs." Though the data is not captured in the report, the authors believe anecdotally that of those who have deployed MD5 somewhere in their network, only about 25-30% of the routers in their network are deployed with the authentication enabled. None report using IPsec to protect the routing protocol, and this was a decline from the few that reported doing so in the previous year's report.

From my personal conversations with operators, of those using MD5, almost all report deploying with one single manual key throughout the entire network. These same operators report that the one single key has not been changed since it was originally installed, sometimes five or more years ago. When asked why, particularly for the case of BGP using TCP MD5, the following reasons are often given:

- A. Changing the keys triggers a TCP reset, and thus bounces the links/adjacencies, undermining Service Level Agreements (SLAs).

B.

For external peers, difficulty of coordination with the other organization is an issue. Once they find the correct contact at the other organization (not always so easy), the coordination function is serialized and on a per peer/AS basis. The coordination is very cumbersome and tedious to execute in practice.

C. Keys must be changed at precisely the same time, or at least within 60 seconds (as supported by two major vendors) in order to limit connectivity outage duration. This is incredibly difficult to do, operationally, especially between different organizations.

D. Relatively low priority compared to other operational issues.

E. Lack of staff to implement the changes device by device.

F. There are three use cases for operational peering at play here: peers and interconnection with other operators, Internal BGP and other routing sessions within a single operator, and operator-to-customer-CPE devices. All three have very different properties, and all are reported as cumbersome. One operator reported that the same key is used for all customer premise equipment. The same operator reported that if the customer mandated, a unique key could be created, although the last time this occurred it created such an operational headache that the administrators now usually tell customers that the option doesn't even exist, to avoid the difficulties. These customer-unique keys are never changed, unless the customer demands so.

The main threat at play here is that a terminated employee from such an operator who had access to the one (or few) keys used for authentication in these environments could easily wage an attack -- or offer the keys to others who would wage the attack -- and bring down many of the adjacencies, causing destabilization to the routing system.

Whatever mechanisms we specify need to be easier than the current methods to deploy, and should provide obvious operational efficiency gains along with significantly better security and threat protection. This combination of value may be enough to drive much broader adoption.

5. Address the threats enumerated above in the ["Threats" section \(Threats\)](#) for each routing protocol, along a roadmap. Not all threats may be able to be addressed in the first specification update for any one protocol. Roadmaps will be defined so that

both the security area and the routing area agree on how the threats will be addressed completely over time.

6. Create a re-usable architecture, framework, and guidelines for various IETF working teams who will address these security improvements for various Routing Protocols. The crux of the KARP work is to re-use that framework as much as possible across relevant Routing Protocols. For example, designers should aim to re-use the key management protocol that will be defined for BGP's TCP-AO key establishment for as many other routing protocols as possible. This is but one example.
7. Bridge any gaps between IETF's Routing and Security Areas by recording agreements on work items, roadmaps, and guidance from the Area leads and Internet Architecture Board (IAB, www.iab.org).

1.5. Non-Goals

[TOC](#)

The following two goals are considered out-of-scope for this effort:

- o Privacy of the packets on the wire, at this point in time. Once this roadmap is realized, we may revisit work on privacy.
- o Message content security. This work is being addressed in other IETF efforts, like SIDR.

1.6. Audience

[TOC](#)

The audience for this roadmap includes:

- o **Routing Area working group chairs and participants** - These people are charged with updates to the Routing Protocol specifications. Any and all cryptographic authentication work on these specifications will occur in Routing Area working groups, with close partnership with the Security Area. Co-advisors from Security Area may often be named for these partnership efforts.
- o **Security Area reviewers of routing area documents** - These people are delegated by the Security Area Directors to perform reviews on routing protocol specifications as they pass through working group last call or IESG review. They will pay particular

attention to the use of cryptographic authentication and corresponding security mechanisms for the routing protocols. They will ensure that incremental security improvements are being made, in line with this roadmap.

- o **Security Area engineers** - These people partner with routing area authors/designers on the security mechanisms in routing protocol specifications. Some of these security area engineers will be assigned by the Security Area Directors, while others will be interested parties in the relevant working groups.
- o **Operators** - The operators are a key audience for this work, as the work is considered to have succeeded if the operators deploy the technology, presumably due to a perception of significantly improved security value coupled with relative similarity to deployment complexity and cost. Conversely, the work will be considered a failure if the operators do not care to deploy it, either due to lack of value or perceived (or real) over-complexity of operations. And as such, the GROW and OPSEC WGs should be kept squarely in the loop as well.

2. Threats

[TOC](#)

In RFC4949[\[RFC4949\]](#) (Shirey, R., "Internet Security Glossary, Version 2," August 2007.), a threat is defined as a potential for violation of security, which exists when there is a circumstance, capability, action, or event that could breach security and cause harm. This section defines the threats that are in scope for this roadmap, and those that are explicitly out of scope. This document leverages the "Generic Threats to Routing Protocols" model, [RFC 4593 \(Barbir, A., Murphy, S., and Y. Yang, "Generic Threats to Routing Protocols," October 2006.\)](#) [RFC4593], capitalizes terms from that document, and offers a terse definition of those terms. (More thorough description of routing protocol threats sources, motivations, consequences and actions can be found in [RFC 4593 \(Barbir, A., Murphy, S., and Y. Yang, "Generic Threats to Routing Protocols," October 2006.\)](#) [RFC4593] itself). The threat listings below expand upon these threat definitions.

2.1. Threats In Scope

[TOC](#)

The threats that will be addressed in this roadmap are those from OUTSIDERS, attackers that may reside anywhere in the Internet, have the ability to send IP traffic to the router, may be able to observe the

router's replies, and may even control the path for a legitimate peer's traffic. These are not legitimate participants in the routing protocol. Message authentication and integrity protection specifically aims to identify messages originating from OUTSIDERS.

The concept of OUTSIDERS can be further refined to include attackers who are terminated employees, and those sitting on-path.

- o On-Path - attackers with control of a network resource or a tap along the path of packets between two routers. An on-path outsider can attempt a man-in-the-middle attack, in addition to several other attack classes. A man-in-the-middle (MitM) attack occurs when an attacker who has access to packets flowing between two peers tampers with those packets in such a way that both peers think they are talking to each other directly, when in fact they are actually talking to the attacker only. Protocols conforming to this roadmap will use cryptographic mechanisms to prevent a man-in-the-middle attacker from situating himself undetected.
- o Terminated Employees - in this context, those who had access router configuration that included keys or keying material like pre-shared keys used in securing the routing protocol. Using this material, the attacker could send properly MAC'd spoofed packets appearing to come from router A to router B, and thus impersonate an authorized peer. The attacker could then send false traffic that changes the network behavior from its operator's design. The goal of addressing this source specifically is to call out the case where new keys or keying material becomes necessary very quickly, with little operational expense, upon the termination of such an employee. This grouping could also refer to any attacker who somehow managed to gain access to keying material, and said access had been detected by the operators such that the operators have an opportunity to move to new keys in order to prevent an attack.

These ATTACK ACTIONS are in scope for this roadmap:

- o SPOOFING - when an unauthorized device assumes the identity of an authorized one. Spoofing can be used, for example, to inject malicious routing information that causes the disruption of network services. Spoofing can also be used to cause a neighbor relationship to form that subsequently denies the formation of the relationship with the legitimate router.
- o FALSIFICATION - an action whereby an attacker sends false routing information. To falsify the routing information, an attacker has to be either the originator or a forwarder of the routing information. Falsification may occur by an ORIGINATOR, or a FORWARDER, and may involve OVERCLAIMING, MISCLAIMING, or

MISTATEMENT of network resource reachability. We must be careful to remember that in this work we are only targeting falsification from outsiders as may occur from tampering with packets in flight. Falsification from BYZANTINES (see the [Threats Out of Scope section \(Threats Out of Scope\)](#) below) are not addressed by the KARP effort.

- o INTERFERENCE - when an attacker inhibits the exchanges by legitimate routers. The types of interference addressed by this work include:
 - *ADDING NOISE
 - *REPLAYING OUT-DATED PACKETS
 - *INSERTING MESSAGES
 - *CORRUPTING MESSAGES
 - *BREAKING SYNCHRONIZATION
 - *Changing message content
- o DoS attacks on transport sub-systems - This includes any other DoS attacks specifically based on the above attack types. This is when an attacker sends spoofed packets aimed at halting or preventing the underlying protocol over which the routing protocol runs, for example halting a BGP session by sending a TCP FIN or RST packet. Since this attack depends on spoofing, operators are encouraged to deploy
- o DoS attacks using the authentication mechanism - This includes an attacker sending packets which confuse or overwhelm a security mechanism itself. An example is initiating an overwhelming load of spoofed authenticated route messages so that the receiver needs to process the MAC check, only to discard the packet, sending CPU levels rising. Another example is when an attacker sends an overwhelming load of keying protocol initiations from bogus sources. All other possible DoS attacks are out of scope (see next section).
- o Brute Force Attacks Against Password/Keys - This includes either online or offline attacks where attempts are made repeatedly using different keys/passwords until a match is found. While it is impossible to make brute force attacks on keys completely unsuccessful, proper design can make such attacks much harder to succeed. For example, the key length should be sufficiently long so that covering the entire space of possible keys is improbable using computational power expected to be available 10 years out or more. Also, frequently changing the keys may render useless a

successful guess some time in the future, as those keys may no longer be in use.

2.2. Threats Out of Scope

[TOC](#)

Threats from BYZANTINE sources -- faulty, misconfigured, or subverted routers, i.e., legitimate participants in the routing protocol -- are out of scope for this roadmap. Any of the attacks described in the above [section \(Threats In Scope\)](#) that may be levied by a BYZANTINE source are therefore also out of scope. In addition, these other attack actions are out of scope for this work:

- *SNIFFING - passive observation of route message contents in flight
 - *FALSIFICATION by BYZANTINE sources - unauthorized message content by a legitimate authorized source.
 - *INTERFERENCE due to:
 - NOT FORWARDING PACKETS - cannot be prevented with cryptographic authentication
 - DELAYING MESSAGES - cannot be prevented with cryptographic authentication
 - DENIAL OF RECEIPT - cannot be prevented with cryptographic authentication
 - UNAUTHORIZED MESSAGE CONTENT - the work of the IETF's SIDR working group (<http://www.ietf.org/html.charters/sidr-charter.html>).
 - Any other type of DoS attack. For example, a flood of traffic that fills the link ahead of the router, so that the router is rendered unusable and unreachable by valid packets is NOT an attack that this work will address. Many other such examples could be contrived.
-

3. Categorizing Routing Protocols

[TOC](#)

For the purpose of this security roadmap definition, we will categorize the routing protocols into groups and have design teams focus on the

specification work within those groupings. It is believed that the groupings will have like requirements for their authentication mechanisms, and that reuse of authentication mechanisms will be greatest within these grouping. The work items placed on the roadmap will be defined and assigned based on these categorizations. It is also hoped that, down the road in the Phase 2 work, we can create one KMP per category (if not for several categories) so that the work can be easily leveraged by the various Routing Protocol teams. KMPs are useful for allowing simple, automated updates of the traffic keys used in a base protocol. KMPs replace the need for humans, or OSS routines, to periodically replace keys on running systems. It also removes the need for a chain of manual keys to be chosen or configured. When configured properly, a KMP will enforce the key freshness policy of two peers by keeping track of the key lifetime and negotiating a new key at the defined interval.

3.1. Category: Messaging Transaction Type

[TOC](#)

The first categorization defines four types of messaging transactions used on the wire by the base Routing Protocol. They are:

One-to-One One peer router directly and intentionally delivers a route update specifically to one other peer router. Examples are BGP and LDP. Point-to-point modes of both IS-IS and OSPF, when sent over both traditional point-to-point links and when using multi-access layers, may both also fall into this category. [question to reviewers: Should we list all protocols into these categories right here, or just give a few examples?]

One-to-Many A router peers with multiple other routers on a single network segment -- i.e. on link local -- such that it creates and sends one route update message which is intended for consumption by multiple peers. Examples would be OSPF and IS-IS in their broadcast, non-point-to-point modes.

Client-Server A client-server routing protocol is one where one router initiates a request for route information from another router, who then formulates a response to that request, and replies with the requested data. Examples are a BGP Route Reflector and [???? Are there other examples? Is this the right example? Do discovery protocols fall under this category?].

Multicast Multicast protocols have unique security properties because of the fact that they are inherently group-based protocols and thus have group keying requirements at the routing level where link-local routing messages are multicasted. Also, at least in the case of PIM-SM, some messages are sent unicast

to a given peer(s), as is the case with router-close-to-sender and the "Rendezvous Point". Some work for application layer message security has been done in the Multicast Security working group (MSEC, <http://www.ietf.org/html.charters/msec-charter.html>) and may be helpful to review, but is not directly applicable.

[author's note: I think the above definitions need clean up. Routing area folks, especially ADs, PLEASE suggest new text.]

3.2. Category: Peer vs. Group Keying

[TOC](#)

The second axis of categorization groups protocols by the keying mechanism that will be necessary for distributing session keys to the actual Routing Protocol transports. They are:

Peer keying One router sends the keying messages directly and only to one other router, such that a one-to-one, unique keying security association (SA) is established between the two routers

Group Keying One router creates and distributes a single keying message to multiple peers. In this case an group SA will be established and used between multiple peers simultaneously. Group keying exists for protocols like [OSPF \(Moy, J., "OSPF Version 2," April 1998.\)](#) [RFC2328] , and also for multicast protocols like [PIM-SM \(Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode \(PIM-SM\): Protocol Specification \(Revised\)," August 2006.\)](#) [RFC4601].

3.3. Category: Update vs. Discovery Protocol

[TOC](#)

The third category group considers protocols by the contents of the messages being exchanged in the Routing Protocol. They are:

Updates Messages that carry route advertisements or update information from peer to peer

Discovery Messages sent as part of a policy, peer, or service discovery process. These messages are normally exchanged prior to any adjacency being formed, and before any updates are sent. For example, end-point discovery mechanisms are common in L2VPN and L3VPN solutions.

[QUESTION TO REVIEWERS: is this really just what's described in 3.1 as "Client-Server" and/or "One-to-One"? Is there really such a different in discovery protocols that they need their own category to figure out how to authenticate them? Can someone provide a few examples?

3.4. Security Characterization Vectors

[TOC](#)

A few more considerations must be made about the protocol and its use when initially categorizing the protocol and scoping the authentication work.

3.4.1. Internal vs. External Operation

[TOC](#)

The designers must consider whether the protocol is an internal routing protocol or an external one, i.e. Does it primarily run between peers within a single domain of control or between two different domains of control? Some protocols may be used in both cases, internally and externally, and as such various modes of authentication operation may be required for the same protocol. While it is preferred that all routing exchanges run with the utmost security mechanisms enabled in all deployments, this exhortation is greater for those protocols running on inter-domain point-to-point links, and greatest for those on shared access link layers with several different domains interchanging together, because the volume of attackers are greater from the outside. Note however that the consequences of internal attacks maybe no less severe -- in fact they may be quite a bit more severe -- than an external attack. An example of this internal versus external consideration is BGP which has both EBGp and IBGP modes. Another example is a multicast protocol where the neighbors are sometimes within a domain of control and sometimes at an inter-domain exchange point. In the case of PIM-SM running on an internal multi-access link, It would be acceptable to give up some security to get some convenience by using a group key between the peers on the link. On the other hand, in the case of PIM-SM running over a multi-access link at a public exchange point, operators may favor security over convenience by using unique pair-wise keys for every peer. Designers must consider both modes of operation and ensure the authentication mechanisms fit both. Operators are encouraged to run cryptographic authentication on all their adjacencies, but to work from the outside in, i.e. The EBGp links are a higher priority than the IBGP links because they are externally facing, and, as a result, more likely to be targeted in an attack.

3.4.2. Unique versus Shared Keys

[TOC](#)

This section discusses security considerations regarding when it is appropriate to use the same authentication key inputs for multiple peers and when it is not. This is largely a debate of convenience versus security. It is often the case that the best secured mechanism is also the least convenient mechanism. For example, an air gap between a host and the network absolutely prevents remote attacks on the host, but having to copy and carry files using the "sneaker net" is quite inconvenient and unscalable.

Operators have erred on the side of convenience when it comes to securing routing protocols with cryptographic authentication. Many do not use it at all. Some use it only on external links, but not on internal links. Those that do use it often use the same key for all peers across their entire network. It is common to see the same key in use for years, and that being the same key that was entered when authentication was originally configured, or the routing gear deployed. The goal for designers is to create authentication mechanisms that are easy for the operators to deploy and manage, and still use unique keys between peers (or small groups on multi-access links), and within between sessions. Operators have the impression that they NEED one key shared across the network, when in fact they do not. What they need is the relative convenience they experience from deploying cryptographic authentication with one (or few) key, compared to the inconvenience they would experience if they deployed the same authentication mechanism using unique pair-wise keys. An example is BGP Route Reflectors. Here operators often use the same authentication key between each client and the route reflector. The roadmaps defined from this guidance document will allow for unique keys to be used between each client and the peer, without sacrificing much convenience. Designers should strive to deliver peer-wise unique keying mechanisms with similar ease-of-deployment properties as today's one-key method. Operators must understand the consequences of using the same keys across many peers. Unique keys are more secure than shared keys because they reduce both the attack target size and the attack consequence size. In this context, the attack target size represents the number of unique routing exchanges across a network that an attacker may be able to observe in order to gain security association credentials, i.e. crack the keys. If a shared key is used across the entire internal domain of control, then the attack target size is very large. The larger the attack target, the easier it is for the attacker to gain access to analysis data, and greater the volume of analysis data he can access in a given time frame, both of which make his job easier. Using the same key across the network makes the attack vulnerability surface more penetrable than unique keys. Consider also the attack consequence size, the amount of routing adjacencies that can be negatively affected once a breach has occurred, i.e. once the keys have been acquired by the attacker. Again, if a shared key is used across the internal

domain, then the consequence size is the whole network. Ideally, unique key pairs would be used for each adjacency.

In some cases designers may need to use shared keys in order to solve the given problem space. For example, a multicast packet is sent once but then observed and consumed by several routing neighbors. If unique keys were used per neighbor, the benefit of multicast would be erased because the casting peer would have to create a different announcement packet/stream for each listening peer. Though this may be desired and acceptable in some small amount of use cases, it is not the norm. Shared group keys are an acceptable solution here, and much work has been done already in this area (see MSEC working group).

3.4.3. Out-of-Band vs. In-line Key Management

[TOC](#)

This section discusses the security and use case considerations for keys placed on devices through out-of-band configurations versus through one routing peer-to-peer key management protocol exchanges. Note, when we say here "Peer-to-Peer KMP" we do not mean in-band to the Routing Protocol. Instead, we mean that the exchange occurs in-line, over IP, between the two routing peers directly. In in-line KMP the peers themselves handle the key and security association negotiation between themselves directly, whereas in an out-of-band system the keys are placed onto the device through some other configuration or management method or interface.

An example of an out-of-band mechanism could be an administrator who makes a remote management connection (e.g. using SSH) to a router and manually enters the keying information -- like the algorithm, the key(s), the lifetimes, etc. Another example could be an OSS system which inputs the same information via a script over an SSH connection, or by pushing configuration through some other management connection, standard (Netconf-based) or proprietary.

The drawbacks of an out-of-band mechanism include: lack of scalability, complexity and speed of changing if a breach is suspected. For example, if an employee who had access to keys was terminated, or if a machine holding those keys was believed compromised, then the system would be considered insecure and vulnerable until new keys were defined by a human. Those keys then need to be placed into the OSS system, manually, and the OSS system then needs to push the change -- often during a very limited change window -- into the relevant devices. If there are multiple organizations involved in these connections, this process is greatly complicated.

The benefits of out-of-band mechanism is that once the new keys/parameters are set in OSS system they can be pushed automatically to all devices within the OSS's domain of control. Operators have mechanisms in place for this already. In small environments with few routers, a manual system is not difficult to employ.

We further define an in-line key exchange as using cryptographically protected identity verification, session key negotiation, and security association parameter negotiation between the two routing peers. The KMP between the two peers may also include the negotiation of parameters, like algorithms, cryptographic inputs (e.g. initialization vectors), key life-times, etc.

The benefits an in-line KMP are several. An in-line KMP results in key(s) that are privately generated, and not recorded permanently anywhere. Since the traffic keys used in a particular connection are not a fixed part of a device configuration no steal-able data exists anywhere else in the operator's systems which can be stolen, e.g. in the case of a terminated or turned employee. If a server or other data store is stolen or compromised, the thieves gain no access to current traffic keys. They may gain access to key derivation material, like a PSK, but not current traffic keys in use. In this example, these PSKs can be updated into the device configurations (either manually or through an OSS) without bouncing or impacting the existing session at all. In the case of using raw asymmetric keys or certificates, instead of PSKs, the data theft would likely not even result in any compromise, as the key pairs would have been generated on the routers, and never leave those routers. In such a case no changes are needed on the routers; the connections will continue to be secure, uncompromised. Additionally, with a KMP regular re-keys operations occur without any operator involvement or oversight. This keeps keys fresh.

The drawbacks to using a KMP are few. First, a KMP requires more cryptographic processing for the router at the very beginning of a connection. This will add some minor start-up time to connection establishment versus a purely manual key approach. Once a connection with traffic keys have been established via a KMP, the performance is the same in the KMP and the out-of-band case. KMPs also add another layer of protocol and configuration complexity which can fail or be misconfigured. This was more of an issue when these KMPs were first deployed, but less so as these implementations and operational experience with them has matured.

The desired end goal is in-line KMPs.

4. The Roadmap

[TOC](#)

4.1. Work Phases on any Particular Protocol

[TOC](#)

The desired endstate for the KARP work contains several items. First, the people desiring to deploy securely authenticated and integrity validated packets between routing peers have the tools specified,

implemented and shipping in order to deploy. These tools should be fairly simple to implement, and not more complex than the security mechanisms to which the operators are already accustomed. (Examples of security mechanisms to which router operators are accustomed include: the use of asymmetric keys for authentication in SSH for router configuration, the use of pre-shared keys (PSKs) in TCP MD5 for BGP protection, the use of self-signed certificates for HTTPS access to device Web-based user interfaces, the use of strongly constructed passwords and/or identity tokens for user identification when logging into routers and management systems.) While the tools that we intend to specify may not be able to stop a deployment from using "foobar" as an input key for every device across their entire routing domain, we intend to make a solid, modern security system that is not too much more difficult than that. In other words, simplicity and deployability are keys to success. The Routing Protocols will specify modern cryptographic algorithms and security mechanisms. Routing peers will be able to employ unique, pair-wise keys per peering instance, with reasonable key lifetimes, and updating those keys on a somewhat regular basis will be operationally easy, causing no service interruption. Achieving the above described end-state using manual keys may only be pragmatic in very small deployments. In larger deployments, this end state will be much more operationally difficult to reach with only manual keys. Thus, there will be a need for key life cycle management, in the form of a key management protocol, or KMP. We expect that the two forms, manual key usage and KMP usage, will co-exist in the real world. For example, a provider's edge router at a public exchange peering point will want to use a KMP for ensuring unique and fresh keys with external peers, while a manual key may be used between a provider's access edge router and each of the same provider's customer premise routers with which it peers. In accordance with the desired end state just described, we define two main work phases for each Routing Protocol:

1. Enhance the Routing Protocol's current authentication mechanism. This work involves enhancing a Routing Protocol's current security mechanisms in order to achieve a consistent, modern level of security functionality within its existing keying framework. It is understood and accepted that the existing keying frameworks are largely based on manual keys. Since many operators have already built operational support systems (OSS) around these manual key implementations, there is some automation available for an operator to leverage in that way, if the underlying mechanisms are themselves secure. In this phase, we explicitly exclude embedding or creating a KMP. A list of the requirements for Phase 1 work are below in the section ["Requirements for Phase](#)

[1 Routing Protocols' Security Updates \(Requirements for Phase 1 Routing Protocols' Security Update\).](#)

2. Develop an automated keying framework. The second phase will focus on the development of an automated keying framework to facilitate unique pair-wise (or perhaps group-wise, where applicable) keys per peering instance. This involves the use of a KMP. A KMP is helpful because it negotiates unique, pair wise, random keys without administrator involvement. It also negotiates several of the SA parameters required for the secure connection, including key life times. It keeps track of those lifetimes using counters, and negotiates new keys and parameters before they expire, again, without administrator interaction. Additionally, in the event of a breach, changing the KMP key will immediately cause a rekey to occur for the Traffic Key, and those new Traffic Keys will be installed and used in the current connection. In summary, a KMP provides a protected channel between the peers through which they can negotiate and pass important data required to exchange proof of key identifiers, derive Traffic Keys, determine re-keying, synchronize their keying state, signal various keying events, notify with error messages, etc. To address brute force attacks [\[RFC3562\] \(Leech, M., "Key Management Considerations for the TCP MD5 Signature Option," July 2003.\)](#) recommends a key management practice to minimize the possibility of successful attack-- frequent key rotation, limited key sharing, key length restrictions, etc. Advances in computational power due to Moore's law are making that management burden untenable-- keys must be of a size and composition that makes configuration and maintenance difficult or keys must be rotated with an unreasonable frequency. A KMP will help immensely with this growing problem.

The framework for any one Routing Protocol will fall under, and be able to leverage, the generic framework described below in section [Section 4.3 \(Common Framework\)](#).

4.2. Requirements for Phase 1 Routing Protocols' Security Update

[TOC](#)

Here is a proposed list of requirements that SHOULD be addressed by Phase 1 (according to "1." above) security updates to Routing Protocols [to be reviewed after -01 is released]:

1. Clear definitions of which elements of the transmission (frame, packet, segment, etc.) are protected by the authentication mechanism

2. Strong algorithms, and defined and accepted by the security community, MUST be specified. The option should use algorithms considered accepted by the security community, which are considered appropriately safe. The use of non-standard or unpublished algorithms SHOULD BE avoided.
3. Algorithm agility for the cryptographic algorithms used in the authentication MUST be specified, i.e. more than one algorithm MUST be specified and it MUST be clear how new algorithms MAY be specified and used within the protocol. This requirement exists in case one algorithm gets broken suddenly. Research to identify weakness in algorithms is constant. Breaking a cipher isn't a matter of if, but when it will occur. It's highly unlikely that two different algorithms will be broken simultaneously. So, if two are supported, and one gets broken, we can use the other until we get a new one in place. Having the ability within the protocol specification to support such an event, having algorithm agility, is essential. Mandating two algorithms provides both a redundancy, and a mechanism for enacting that redundancy when needed.
4. Secure use of simple PSKs, offering both operational convenience as well as building something of a fence around stupidity, MUST be specified.
5. Inter-connection replay protection. Packets captured from one connection MUST NOT be able to be re-sent and accepted during a later connection.
6. Intra-connection replay protection. Packets captured during a connection MUST NOT be able to be re-sent and accepted during that same connection, to deal with long-lived connections.
7. A change of security parameters REQUIRES, and even forces, a change of session traffic keys
8. Intra-connection re-keying which occurs without a break or interruption to the current peering session, and, if possible, without data loss, MUST be specified.
9. Efficient re-keying SHOULD be provided. The specification SHOULD support rekeying during a connection without the need to expend undue computational resources. In particular, the specification SHOULD avoid the need to try/compute multiple keys on a given packet.
10. Prevent DoS attacks as those described as in-scope in the threats section [Section 2.1 \(Threats In Scope\)](#) above.

11. Default mechanisms and algorithms specified and defined as REQUIRED for all implementations
12. Manual keying MUST be supported.
13. Convergence times of the Routing Protocols SHOULD NOT be materially affected. Materially here is defined as anything greater than a 5% convergence time increase. Note that convergence is different than boot time. Also note that convergence time has a lot to do with the speed of processors used on individual routing peers, and this increases by Moore's law over time. Therefore, this requirement should be considered only in terms of total number of messages that must be exchanged, and less for the computational intensity of processing any one message.
14. The changes or addition of security mechanisms SHOULD NOT cause a refresh of route updates or cause additional route updates to be generated
15. Architecture of the specification MUST consider and allow for future use of a KMP.

4.3. Common Framework

[TOC](#)

Each of the categories of routing protocols above will require unique designs for authenticating and integrity checking their protocols. However, a single underlying framework for delivering automatic keying to those solutions will be pursued. Providing such a single framework will significantly reduce the complexity of each step of the overall roadmap. For example, if each Routing Protocol needed to define it's own key management protocol this would balloon the total amount of different sockets that are needed to be opened and processes that are needed to be simultaneously running on an implementation. It would also significantly increase the run-time complexity and memory requirements of such systems running multiple Routing Protocols, causing perhaps slower performance of such systems. However, if we can land on a very small set (perhaps one or two) of automatic key management protocols, KMPs, that the various Routing Protocols can use, then we can reduce this implementation and run-time complexity. We can also decrease the total amount of time implementers need to deliver the KMPs for the Routing Protocols that will provide better threat protection. The components for the framework are listed here, and described below:

*Routing Protocol security mechanism

*KMP

*KeyStore

*Traffic Key

*RoutingProtocol-to-KMP API

*RoutingProtocol-to-KeyStore API

*KMP-to-KeyStore API

*Common Routing Protocol mechanisms

*Identifiers

*Proof of identity

*Profiles

The framework is modularized for how keys and security association (SA) parameters generally get passed from a KMP to a transport protocol. It contains three main blocks and APIs.

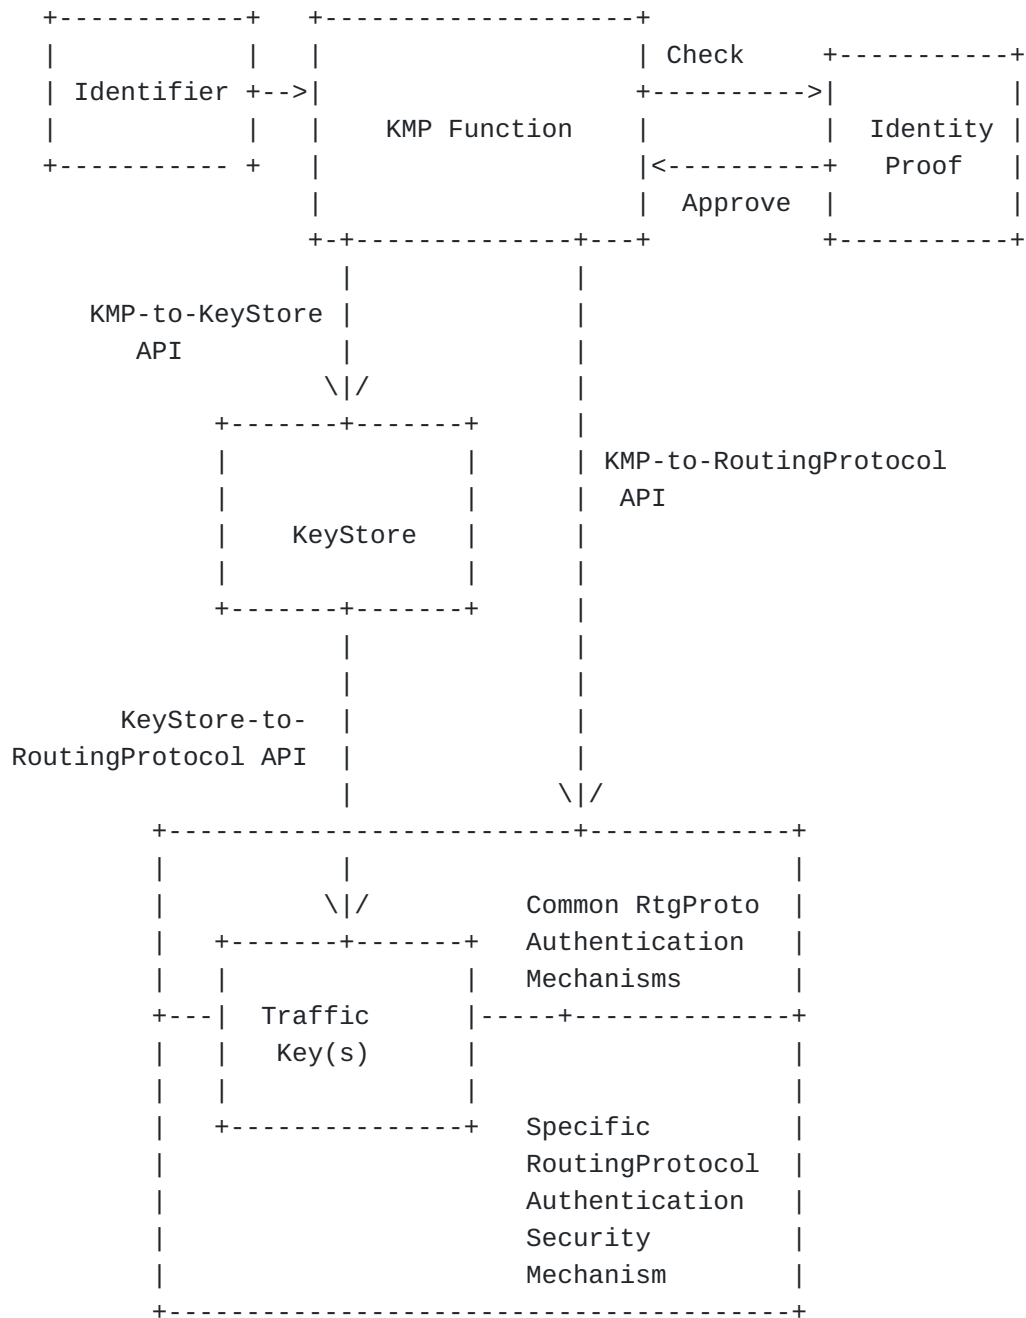


Figure 1: Automatic Key Management Framework

Each element of the framework is described here:

- o **Routing Protocol** - Routing protocol security mechanism - In each case, the Routing Protocol will contain a mechanism for using session keys in their security option. When the Routing Protocol uses a transport substrate, e.g. the way BGP, LDP and MSDP use

TCP, then this applies to the security mechanism the includes that substrate.

- o **KeyStore** - Each implementation will also contain a protocol independent mechanism for storing keys, called KeyStore. The KeyStore will have multiple different logical containers, one container for each session key that any given Routing Protocol will need. Keys stored here may be a Peer Key or a Traffic Key. There may also be associated parameters as required by the SA for any given Routing Protocol.
- o **Peer Key** A key used between peers from which a traffic key is derived. An example is a PSK.
- o **Traffic Key** The actual key used on each packet of a message. This key may be derived from the key existing in the KeyStore. This will depend on whether the key in KeyStore was a manual PSK for the peers, or whether a connection-aware KMP created the key. Further, it will be connection specific, so as to provide inter- and intra-connection replay protection.
- o **RoutingProtocol-KeyStore API** - There will be an API for Routing Protocol to retrieve (or receive; it could be a push or a pull) the keys from the KeyStore. This will enable implementers to reuse the same API calls for all their Routing Protocols. The API will necessarily include facility to retrieve other SA parameters required for the construction of the Routing Protocol's packets, like key IDs or key lifetimes, etc.
- o **KMP** - There will be an automated key management protocol, KMP. This KMP will run between the peers. The KMP serves as a protected channel between the peers, through which they can negotiate and pass important data required to exchange proof of key identifiers, derive session keys, determine re-keying, synchronize their keying state, signal various keying events, notify with error messages, etc. As an analogy, in the IPsec protocol ([RFC4301 \(Kent, S. and K. Seo, "Security Architecture for the Internet Protocol," December 2005.\)](#) [RFC4301], [RFC4303 \(Kent, S., "IP Encapsulating Security Payload \(ESP\)," December 2005.\)](#) [RFC4303] and [RFC4306 \(Kaufman, C., "Internet Key Exchange \(IKEv2\) Protocol," December 2005.\)](#) [RFC4306]) IKEv2 is the KMP that runs between the two peers, while AH and ESP are two different base protocols that take session keys from IKEv2 and use them in their transmissions. In the analogy, the Routing Protocol, say BGP and LDP, are analogous to ESP and AH, while the KMP is analogous to IKEv2 itself.
- o **RoutingProtocol-KMP API** - There will be an API for the Routing Protocol to request a session key of the KMP, and be notified

when the keys are available for it. The API will also contain a mechanism for the KMP to notify the Routing Protocol that there are new keys that it must now use, even if it didn't request those keys. The API will also include a mechanism for the KMP to receive requests for session keys and other parameters from the routing protocol. The KMP will also be aware of the various Routing Protocols and each of their unique parameters that need to be negotiated and returned.

- o **KMP-KeyStore API** - There will be an API for the KMP to place keys and parameters into the KeyStore after their negotiation and derivation with the other peer. This will enable the implementers to reuse the same calls for multiple KMPs that may be needed to address the various categories of Routing Protocols as described in the section defining [categories \(Categorizing Routing Protocols\)](#).

[after writing this all up, I'm not sure we really need the key_store in the middle. As long as we standardize fully all the calls needed from any Routing Protocol to any KMP, then there can be a generic hand-down function from the KMP to the Routing Protocol when the key and parameters are ready. Let's sleep on it.]

[will need state machines and function calls for these APIs, as one of the work items. In essence, there is a need for a core team to develop the APIs out completely in order for the Routing Protocol teams to use them. Need to get this team going asap.]

- o **Identifiers** - A KMP is fed by identities. The identities are text strings used by the peers to indicate to each other that each are known to the other, and authorized to establish connections. Those identities must be represented in some standard string format, e.g. an IP address -- either v4 or v6, an FQDN, an RFC 822 email address, a Common Name [RFC PKI], etc. Note that even though routers do not normally have email addresses, one could use an RFC 822 email address string as a formatted identifier for a router. They would do so simply by putting the router's reference number or name-code as the "NAME" part of the address, left of the "@" symbol. They would then place some locational context in the "DOMAIN" part of the string, right of the "@" symbol. An example would be "rtr0210@sf.ca.us.company.com". This document does not suggest this string value at all. Instead, the concept is used only to clarify that the type of string employed does not matter. It also does not matter what specific text you chose to place in that string type. It only matters that the type of string -- and its format -- must be agreed upon by the two endpoints. Further, the string can be used as an identifier in this context, even if the string is not actually provisioned in its source domain. For example, the email address "rtr0210@sf.ca.us.company.com" may not actually exist as an email

address in that domain, but that string of characters may still be used as an identifier type(s) in the routing protocol security context. What is important is that the community decide on a small but flexible set of Identifiers they will all support, and that they decide on the exact format of those string. The formats that will be used must be standardized and must be sensible for the routing infrastructure.

- o Identity Proof** - Once the form of identity is decided, then there must be a cryptographic proof of that identity, that the peer really is who they assert themselves to be. Proof of identity can be arranged between the peers in a few ways, for example pre-shared keys, raw assymetric keys, or a more user-friendly representation of assymetric keys, like a certificate. Certificates can be used in a way requiring no additional supporting systems -- e.g. public keys for each peer can be maintained locally for verification upon contact. Certificate management can be made more simple and scalable with the use minor additional supporting systems, as is the case with self-signed certificates and a flat file list of "approved thumbprints". Self-signed certificates will have somewhat lower security properties than Certificate Authority signed certificates [RFC Certs]. The use of these different identity proofs vary in ease of deployment, ease of ongoing management, startup effort, ongoing effort and management, security strength, and consequences from loss of secrets from one part of the system to the rest of the system. For example, they differ in resistance to a security breach, and the effort required to remediate the whole system in the event of such a breach. The point here is that there are options, many of which are quite simple to employ and deploy.
- o Profiles** - Once the KMP, Identifiers and Proofs mechanisms are converged upon, they must be clearly profiled for each Routing Protocol, so that implementors and deployers alike understand the different pieces of the solution, and can have similar configurations and interoperability across multiple vendors' devices, so as to reduce management difficulty. The profiles SHOULD also provide guidance on when to use which various combinations of options. This will, again, simplify use and interoperability.

In addition to other business, administrative, and operational terms they must already exchange prior to forming first adjacencies, it is assumed that two parties deploying message authentication on their routing protocol will also need to decide upon acceptable security parameters for the connection. This will include the form and content of the identity each use to represent the other. It will also include the type of keys to be used, e.g. PSK, raw assymetric keys,

certificate. And it will include the acceptable cryptographic algorithms, or algorithm suite. This agreement is necessary in order for each to properly configure the connection on their respective devices. The manner in which they agree upon and exchange this policy information is normally via phone call or written exchange, and is outside the scope of the KARP effort, but assumed to have occurred. We take as a given that each party knows the identity types and values, key types and values, and acceptable cryptographic algorithms for both their own device and the peer that form the security policy for configuration on their device.

Common Mechanisms - In as much as they exist, the framework will capture mechanisms that can be used commonly not only within a particular category of Routing Protocol and Routing Protocol to KMP, but also between Routing Protocol categories. Again, the goal here is simplifying the implementations and runtime code and resource requirements. There is also a goal here of favoring well vetted, reviewed, operationally proven security mechanisms over newly brewed mechanisms that are less well tried in the wild.

4.4. Work Items Per Routing Protocol

[TOC](#)

Each Routing Protocol will have a team (the [Routing_Protocol]-KARP team) working on incrementally improving their Routing Protocol's security. These teams will have the following main work items:

PHASE 1:

Characterize the RP Assess the Routing Protocol to see what authentication mechanisms it has today. Does it need significant improvement to its existing mechanisms or not? This will include determining if modern, strong security algorithms and parameters are present.

Define Optimal State List the requirements for the Routing Protocol's session key usage and format to contain to modern, strong security algorithms and mechanisms, per the [Requirements \(Requirements for Phase 1 Routing Protocols' Security Update\)](#) section above. The goal here is to determine what is needed for the Routing Protocol alone to be used securely with at least manual keys.

Gap Analysis Enumerate the requirements for this protocol to move from its current security state, the first bullet, to its optimal state, as listed just above.

Transition and Deployment Considerations Document the operational transition plan for moving from the old to the new security mechanism. Will adjacencies need to bounce? What new elements/

servers/services in the infrastructure will be required? What is an example work flow that an operator will take? The best possible case is if the adjacency does not break, but this may not always be possible.

Define, Assign, Design Create a deliverables list of the design and specification work, with milestones. Define owners. Release a document(s)

PHASE 2:

KMP Analysis Review requirements for KMPs [RFC????]. Identify any nuances for this particular protocol's needs and its use cases for KMP. List the requirements that this Routing Protocol has for being able to be use in conjunctions with a KMP. Define the optimal state.

Gap Analysis Enumerate the requirements for this protocol to move from its current security state to its optimal state.

Define, Assign, Design Create a deliverables list of the design and specification work, with milestones. Define owners. Do the design and document work for a KMP to be able to generate the Routing Protocol's session keys for the packets on the wire. These will be the arguments passed in the API to the KMP in order to bootstrap the session keys to the Routing Protocol.

There will also be a team formed to work on the base framework mechanisms for each of the main categories, i.e. the blocks and API's represented in [figure 1 \(Automatic Key Management Framework\)](#).

4.5. Protocols in Categories

[TOC](#)

This section groups the Routing Protocols into like categories, according to attributes set forth in [Categories Section \(Categorizing Routing Protocols\)](#). Each group will have a design team tasked with improving the security of the Routing Protocol mechanisms and defining the KMP requirements for their group, then rolling both into a roadmap document upon which they will execute.

BGP, LDP and MSDP The Routing Protocol's that fall into the category of the one-to-one peering messages, and will use peer keying protocols, AND are all transmitted over TCP include BGP [RFC 4271 \(Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 \(BGP-4\)," January 2006.\)](#) [RFC4271], LDP ([Andersson, L., Minei, I., and B. Thomas, "LDP Specification," October 2007.](#)) [RFC5036] and MSDP ([Fenner, B. and D. Meyer, "Multicast Source](#)

[Discovery Protocol \(MSDP\),” October 2003.\)](#) [RFC3618]. A team will work on one mechanism to cover these three protocols. Much of the work on the Routing Protocol update for its existing authentication mechanism is already occurring in the TCPM Working Group, on the [TCP-AO \(Touch, J., Mankin, A., and R. Bonica, “The TCP Authentication Option,” March 2010.\)](#)

[I-D.ietf-tcpm-tcp-auth-opt] document, as well as its cryptography-helper document, [TCP-AO-CRYPTO \(Lebovitz, G., “Cryptographic Algorithms, Use and Implementation Requirements for TCP Authentication Option,” March 2009.\)](#) [I-D.ao-crypto]. The exception is the mode where LDP is used directly on the LAN [RFC????]. The work for this may go into the Group keying category (w/ OSPF) mentioned below.

OSPF, ISIS, and RIP The Routing Protocols that fall into the category Group keying with one-to-many peering messages includes [OSPF \(Moy, J., “OSPF Version 2,” April 1998.\)](#) [RFC2328], [ISIS \(Callon, R., “Use of OSI IS-IS for routing in TCP/IP and dual environments,” December 1990.\)](#) [RFC1195] and [RIP \(Malkin, G., “RIP Version 2,” November 1998.\)](#) [RFC2453]. Not surprisingly, all these routing protocols have two other things in common. First, they are run on a combination of the OSI datalink layer 2, and the OSI network layer 3. By this we mean that they have a component of how the routing protocol works which is specified in Layer 2 as well as in Layer 3. Second, they are all internal gateway protocols, or IGP's. The keying mechanisms and use will be much more complicated to define for these than for a one-to-one messaging protocol.

BFD Because it is less of a routing protocol, per se, and more of a peer aliveness detection mechanism, Bidirectional Forwarding Detection (BFD) [RFC????] will have its own team.

RSVP [RFC????], RSVP-TE [RFC????], and PCE These three protocols will be handled together. [what more characterisation should we give here? Routing AD's, provide text pls?]

PIM-SM and PIM-DM Finally, the multicast protocols of [PIM-SM \(Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, “Protocol Independent Multicast - Sparse Mode \(PIM-SM\): Protocol Specification \(Revised\),” August 2006.\)](#) [RFC4601] and [PIM-DM \(Adams, A., Nicholas, J., and W. Siadak, “Protocol Independent Multicast - Dense Mode \(PIM-DM\): Protocol Specification \(Revised\),” January 2005.\)](#) [RFC3973] will be handled together. PIM-SM multicasts routing information (Hello, Join/Prune, Assert) on a link-local basis, using a defined multicast address. In addition, it specifies unicast communication for exchange of information (Register, Register-Stop) between the router closest to a group sender and the "rendezvous point" (RP). The RP is

typically not "on-link" for a particular router. While much work has been done on multicast security for application-layer groups, little has been done to address the problem of managing hundreds or thousands of small one-to-many groups with link-local scope. Such an authentication mechanism should be considered along with the router-to-Rendezvous Point authentication mechanism. The most important issue is ensuring that only the "authorized neighbors" get the keys for (S,G), so that rogue routers cannot participate in the exchanges. Another issue is that some of the communication may occur intra-domain, e.g. the link-local messages in an enterprise, while others for the same (*,G) may occur inter-domain, e.g. the router-to-Rendezvous Point messages may be from one enterprise's router to another. One possible solution proposes a region-wide "master" key server (possibly replicated), and one "local" key server per speaking router. There is no issue with propagating the messages outside the link, because link-local messages, by definition, are not forwarded. This solution is offered only as an example of how work may progress; further discussion should occur in this work team. Specification of a link-local protection mechanism for PIM-SM occurred in [RFC 4601 \(Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode \(PIM-SM\): Protocol Specification \(Revised\)," August 2006.\)](#) [RFC4601], and this work is being updated in [PIM-SM-LINKLOCAL \(Atwood, W., Islam, S., and M. Siami, "Authentication and Confidentiality in PIM-SM Link-local Messages," December 2009.\)](#) [I-D.ietf-pim-sm-linklocal]. However, the KMP part is completely unspecified, and will require work outside the expertise of the PIM working group to accomplish, which is why this roadmap is being created.

These protocols are deemed out-of-scope for this current iteration of the work roadmap. Once all of the protocols listed above have had their work completed, or are clearly within site of completion, then the community will revisit the need and interest for working on these:

*MANET

*FORCES

[need text from routing ADs on why these are out of scope]

4.6. Priorities

[TOC](#)

Resources from both the routing area and the security area will be applied to work on these problem spaces as quickly as possible. Realizing that such resources are far from unlimited, a rank order

priority for addressing the work of incrementally securing these groups of routing protocols is provided:

- *Priority 1 - BGP / LDP / MSDP - almost done with Phase 1 on these, via TCP-AO [\[I-D.ietf-tcpm-tcp-auth-opt\] \(Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option," March 2010.\)](#) .

- *Priority 2 - PIM-SM

- *Priority 3 - OSPF / ISIS / RIP

- *Priority 4 - BFD

- *Priority 5 - RSVP and RSVP-TE

By far the most important group is the Priority 1 group as these are the protocols used on the most public and exposed segments of the networks, at the peering points between operators and between operators and their customers. BFD, as a detection mechanism underlying the Priority 1 protocols is therefore second.

5. Security Considerations

[TOC](#)

As mentioned in the Introduction , RFC4948 identifies additional steps needed to achieve the overall goal of improving the security of the core routing infrastructure. Those include validation of route origin announcements, path validation, cleaning up the IRR databases for accuracy, and operational security practices that prevent routers from being compromised devices. The KARP work is but one step in a necessary system of security improvements.

The security of cryptographic-based systems depends on both the strength of the cryptographic algorithms chosen and the strength of the keys used with those algorithms. The security also depends on the engineering of the protocol used by the system to ensure that there are no non-cryptographic ways to bypass the security of the overall system. Care should also be taken to ensure that the selected key is unpredictable, avoiding any keys known to be weak for the algorithm in use. [\[RFC4086\] \(Eastlake, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security," June 2005.\)](#) contains helpful information on both key generation techniques and cryptographic randomness.

In addition to using a strong key/PSK of appropriate length and randomness, deployers of KARP protocols SHOULD use different keys between different routing peers whenever operationally possible. [RFC3562 \(Leech, M., "Key Management Considerations for the TCP MD5 Signature Option," July 2003.\)](#) [RFC3562] provides some very sound guidance. It was meant specifically for the use of TCP MD5 for BGP, but

it is more or less applicable to Routing Protocol authentication work that would result from KARP. It states three main points: (1) key lengths SHOULD be between 12 and 24 bytes (this will vary depending on the MAC/KDF in use), with larger keys having effectively zero additional computational costs when compared to shorter keys, (2) key sharing SHOULD be limited so that keys aren't shared among multiple BGP peering arrangements, and (3) Keys SHOULD be changed at least every 90 days (this could be longer for stronger MAC algorithms, but it is generally a wise idea).

This is especially true when the Routing Protocol takes a static Traffic Key as opposed to a Traffic Key derived per-connection by a KDF. The burden for doing so is understandable much higher than for using the same static Traffic Key across all peering routers. This is why use of a KMP network-wide increases peer-wise security so greatly, because now each set of peers can enjoy a unique Traffic Key, and if an attacker sitting between two routers learns or guesses the Traffic Key for that connection, she doesn't gain access to all the other connections as well.

However, whenever using manual keys, it is best to design a system where a given PSK will be used in a KDF, mixed with connection specific material, in order to generate session unique -- and therefore peer-wise -- Traffic Keys. Doing so has the following advantages: the Traffic Keys used in the per-message MAC operation are peer-wise unique, it provides inter-connection replay protection, and, if the per-message MAC covers some connection counter, intra-connection replay protection.

Note that in the composition of certain key derivation functions (e.g. KDF_AES_128_CMAC, as used in TCP-AO [\[I-D.ao-crypto\] \(Lebovitz, G., "Cryptographic Algorithms, Use and Implementation Requirements for TCP Authentication Option," March 2009.\)](#)), the pseudorandom function (PRF) used in the KDF may require a key of a certain fixed size as an input. For example, AES_128_CMAC requires a 128 bit (16 byte) key as the seed. However, for convenience to the administrators/deployers, a specification may not want to force the deployer to enter a PSK of exactly 16 bytes. Instead, a specification may call for a sub-key routine that could handle a variable length PSK, one that might be less than 16 bytes (see [\[RFC4615\] \(Song, J., Poovendran, R., Lee, J., and T. Iwata, "The Advanced Encryption Standard-Cipher-based Message Authentication Code-Pseudo-Random Function-128 \(AES-CMAC-PRF-128\) Algorithm for the Internet Key Exchange Protocol \(IKE\)," August 2006.\)](#), section 3, as an example). That sub-key routine would act as a key extractor to derive a second key of exactly the required length, and thus suitable as a seed to the PRF. This does NOT mean that administrators are safe to use weak keys. Administrators are encouraged to follow [\[RFC4086\] \(Eastlake, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security," June 2005.\)](#) as listed above. We simply attempted to "put a fence around stupidity", in as much as possible.

A better option, from a security perspective, is to use some representation of a device-specific asymmetric key pair as the identity proof, as described in [Section 3.4.2 \(Unique versus Shared Keys\)](#).

When it comes time for the KARP WG to design the re-usable model for a KMP, [The Guidelines for Cryptographic Key Management, RFC4107](#)

[\(Bellovin, S. and R. Housley, "Guidelines for Cryptographic Key Management," June 2005.\)](#) [RFC4107] should be will be consulted.

[[QUESTION TO REVIEWERS: it may be worthwhile to pull the last few paragraphs, along with some guidance along the same lines, into section 4, in a new sub-section with a title something like "Security tips for KARP design teams working on Routing Protocol reviews and updates". Or maybe even into its own info document, "Security Guidelines for KARP Design Teams".Thoughts?]]

The mechanisms that will be defined under this roadmap aim to improve the security, better protect against more threats, and provide far greater operational efficiencies than the state of things at the time of this writing. However, none of these changes will improve Internet security unless they are implemented and deployed. Other influences must be brought to bear upon operators and organizations to create incentives for deployment. Such incentives may take the form of PCI-like industry compliance/certifications, well advertised BCPs profiling the use of this roadmap's output, end-user demand or insistence.

6. IANA Considerations

[TOC](#)

This document has no actions for IANA.

7. Acknowledgements

[TOC](#)

The outline for this draft was created from discussions and agreements with Routing AD's Ross Callon and Dave Ward, Security AD's Tim Polk and Pasi Eronen, and IAB members Danny McPherson and Gregory Lebovitz.

Mat Ford and Bill Atwood provided reviews to -00.

Danny McPherson provided an extremely detailed and useful review of -01.

8. Change History (RFC Editor: Delete Before Publishing)

[TOC](#)

[NOTE TO RFC EDITOR: this section for use during I-D stage only. Please remove before publishing as RFC.]

kmart-00-00 original rough rough rough draft for review by routing and security AD's

kmart-00- original submission

- *adds new category = multicast protocols in category section and mentions mcast in group keying category description.
- *add a lot of references where they did not exist before, or where there were only place holders. Still more work needed on this.
- *abstract filled in
- *changed from standards track to informational (this was an oversight in last draft).
- *filled out threats section with detailed descriptions, and linked to RPsec threats RFC
- *made ascii art for the basic KMP framework
- *added section on internal versus external peering and the requirements decisions for them
- *added security characterization section in sect 2, added sections discussing internal vs external protocols, shared vs unique keys, oob vs in-band keying
- *incorporates all D Ward's feedback from his initial skim of the document.

kmart-01 -

- *Updated [framework \(Automatic Key Management Framework\)](#) diagram to include all listed and described elements. Needs review and honing. Gregory Lebovitz (GL).
- *Added comment in [protocols \(Protocols in Categories\)](#) section that much of the BGP/LDP Phase 1 work is already being done in tcp-ao and ao-crypto. GL.
- *Updated Scope making the 2 work phases more clear earlier in the document. GL.
- *Broke [work items \(Work Items Per Routing Protocol\)](#) section into two Phases, 1 for manual key update, and second for KMP work. GL.
- *Re-org'd doc. Brought [Threats \(Threats\)](#) section out into its own top-level section. Did same with [Categorization \(Categorizing Routing Protocols\)](#) section, leaving Roadmap section more focused.

Moved ToDo list and Change History to end of doc, after Acknowledgements. GL.

*added new [sect 2.3 \(Work Phases on any Particular Protocol\)](#) on main roadmap phases. Previous section [Common Framework \(Common Framework\)](#) moved to 2.4. Tim Polk (TP).

*Added Section 2.3.1 [Requirements for Phase 1 Routing Protocols' Security Update \(Requirements for Phase 1 Routing Protocols' Security Update\)](#). This provides a nice starter set of requirements for any work team. GL.

*Filled out text for [Out vs In-band Key Mgmt \(Out-of-Band vs. In-line Key Management\)](#) section, significantly. Changed the term from "in-band" to "in-line".

*Section [Threats \(Threats\)](#) Clarified DoS threats in and out of scope better. We are not preventing all DoS attacks. Just those we can reasonably via authentication. TP.

*Sect [In-band vs Out-of-Band \(Out-of-Band vs. In-line Key Management\)](#) clarified that In-band does not mean in-band to Routing Protocol, but rather over IP between the Routing Protocols, rather than pushed down by some external management entity. TP.

*In [roadmap \(Categorizing Routing Protocols\)](#) section, added "it is also hoped that we can create one kmp per category..." Also explained value of a KMP. TP.

*Added "operators" to [audience \(Audience\)](#) list. Matt Ford (MF).

*Described why BGP (and LDP) security is not deployed very often. Added this [Scope \(Scope\)](#) section, point 4. If mechanisms aren't being deployed, why is that? What, if anything, could be done to improve deployment? Tried to address these. Need references (see To Do list below). MF.

*Added some text to security section to address this from MF: say something here about the limitations of this approach, if any - and refer back to the need for other pieces of the puzzle. May need more work.

*Cleaned up text for multicast part of [Message Type \(Category: Messaging Transaction Type\)](#) section and [Protocols \(Protocols in Categories\)](#) section, clarifying PIM's two message types, mcast and unicast, in both places. Bill Atwood (BA).

*In section [Protocols \(Protocols in Categories\)](#), added references to 4601 and PIM-SM-LINKLOCAL. BA.

*Editorial changes pointed out various folks.

kmart-02 -

*Re-submitted due to expiration. Text did not change. Substantive update coming shortly.

kmart-03 -

*changed "BaseRP" to "Routing Protocol" throughout the doc - man

*filled out the Terminology section

*changed "KMART" to "KARP" in everything but the title, since the -00 deadline had long since passed. Will change the title of the doc to KARP as soon as the window re-opens.

*priorities in sect 4.6 changed. Added PIM-SM. Lowered OSPF and BFD, based on feedback by a few people.

*many edits resulting from Danny McPherson's review.

*added "Brute Force Attacks Against Password/Keys" to Threats [Section 2.1 \(Threats In Scope\)](#) section.

*Significant updates to Security Considerations section

*Added a few references throughout to RFC3562

*4.3 2nd to last P - added a comment to clarify that two parties (or an org) must discuss ahead of time what they want their connections' security properties to be. - dward

*added to 4.4 Phase 1 - New Section: Transition and Deployment Considerations. ea wg must call out the operational transition plan from old to new security. Best if don't bounce link. - dward

*added 3.3 (but not sure if this is right)- endpoint discovery mechanisms? endpoint discovery mechanism (L2VPN, L3VPN, etc). Discovery is much different security properties than passing Routing updates. - dward

*More requirements: Added to 4.2: X - convergence SHOULD not be affected by what we choose; adding security SHOULD not cause a refresh of route updates or cause additional route updates to be generated; adding auth should not be an attack vector itself. AKA, the use of MD5 is so expensive that spoofing BGP packets w/ MD5 causes the control plane to be attacked because CPU went to 100% - dward

*updated stats on MD5 usage, and cited [ISR2008]. - mchpherson

karp-00 -

*changes title from "kmart" to "karp" and the version from "-03" to "00". No other changes.

9. Needs Work in Next Draft (RFC Editor: Delete Before Publishing)

[TOC](#)

[NOTE TO RFC EDITOR: this section for use during I-D stage only. Please remove before publishing as RFC.]

List of stuff that still needs work

*RTG AD's or delegates: clean up the three definitions of route message type categories. Need RTG Area folks input on this.

*More clarity on the work items for those defining and specifying the framework elements and API's themselves.

*RTG AD's or delegates: text justifying RSVP and RSVP-TE and what we think solving that problem may look like

*RTG AD's or delegates: more justification for why MANET and FORCES are out of scope. Need ref for those RFCs.

*Danny McPherson: Get reference for BGP auth usage stats in [Scope \(Scope\)](#) section, item 4.

*security section: pull out security guidance to routing protocol design teams stuff and place into its own section?

*

10. References

[TOC](#)

10.1. Normative References

[TOC](#)

| | |
|-----------|--|
| [RFC2119] | Bradner, S. , " Key words for use in RFCs to Indicate Requirement Levels ," BCP 14, RFC 2119, March 1997 (TXT , HTML , XML). |
| [RFC4593] | |

| | |
|-----------|--|
| | Barbir, A., Murphy, S., and Y. Yang, " Generic Threats to Routing Protocols ," RFC 4593, October 2006 (TXT). |
| [RFC4948] | Andersson, L., Davies, E., and L. Zhang, " Report from the IAB workshop on Unwanted Traffic March 9-10, 2006 ," RFC 4948, August 2007 (TXT). |

10.2. Informative References

[TOC](#)

| | |
|------------------------------|--|
| [I-D.ao-crypto] | Lebovitz, G., " Cryptographic Algorithms, Use and Implementation Requirements for TCP Authentication Option ," March 2009. |
| [I-D.ietf-pim-sm-linklocal] | Atwood, W., Islam, S., and M. Siami, " Authentication and Confidentiality in PIM-SM Link-local Messages ," draft-ietf-pim-sm-linklocal-10 (work in progress), December 2009 (TXT). |
| [I-D.ietf-tcpm-tcp-auth-opt] | Touch, J., Mankin, A., and R. Bonica, " The TCP Authentication Option ," draft-ietf-tcpm-tcp-auth-opt-11 (work in progress), March 2010 (TXT). |
| [ISR2008] | McPherson, D. and C. Labovitz, " Worldwide Infrastructure Security Report ," October 2008. |
| [RFC1195] | Callon, R., " Use of OSI IS-IS for routing in TCP/IP and dual environments ," RFC 1195, December 1990 (TXT , PS). |
| [RFC2328] | Moy, J., " OSPF Version 2 ," STD 54, RFC 2328, April 1998 (TXT , HTML , XML). |
| [RFC2453] | Malkin, G., " RIP Version 2 ," STD 56, RFC 2453, November 1998 (TXT , HTML , XML). |
| [RFC3562] | Leech, M., " Key Management Considerations for the TCP MD5 Signature Option ," RFC 3562, July 2003 (TXT). |
| [RFC3618] | Fenner, B. and D. Meyer, " Multicast Source Discovery Protocol (MSDP) ," RFC 3618, October 2003 (TXT). |
| [RFC3973] | Adams, A., Nicholas, J., and W. Siadak, " Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised) ," RFC 3973, January 2005 (TXT). |
| [RFC4086] | Eastlake, D., Schiller, J., and S. Crocker, " Randomness Requirements for Security ," BCP 106, RFC 4086, June 2005 (TXT). |
| [RFC4107] | Bellovin, S. and R. Housley, " Guidelines for Cryptographic Key Management ," BCP 107, RFC 4107, June 2005 (TXT). |
| [RFC4271] | Rekhter, Y., Li, T., and S. Hares, " A Border Gateway Protocol 4 (BGP-4) ," RFC 4271, January 2006 (TXT). |
| [RFC4301] | Kent, S. and K. Seo, " Security Architecture for the Internet Protocol ," RFC 4301, December 2005 (TXT). |
| [RFC4303] | |

| | |
|-----------|---|
| | Kent, S., " IP Encapsulating Security Payload (ESP) ," RFC 4303, December 2005 (TXT). |
| [RFC4306] | Kaufman, C., " Internet Key Exchange (IKEv2) Protocol ," RFC 4306, December 2005 (TXT). |
| [RFC4601] | Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, " Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised) ," RFC 4601, August 2006 (TXT , PDF). |
| [RFC4615] | Song, J., Poovendran, R., Lee, J., and T. Iwata, " The Advanced Encryption Standard-Cipher-based Message Authentication Code-Pseudo-Random Function-128 (AES-CMAC-PRF-128) Algorithm for the Internet Key Exchange Protocol (IKE) ," RFC 4615, August 2006 (TXT). |
| [RFC4949] | Shirey, R., " Internet Security Glossary, Version 2 ," RFC 4949, August 2007 (TXT). |
| [RFC5036] | Andersson, L., Minei, I., and B. Thomas, " LDP Specification ," RFC 5036, October 2007 (TXT). |
| [RFC5226] | Narten, T. and H. Alvestrand, " Guidelines for Writing an IANA Considerations Section in RFCs ," BCP 26, RFC 5226, May 2008 (TXT). |

Authors' Addresses

[TOC](#)

| | |
|--------|--|
| | Gregory Lebovitz |
| | Juniper Networks, Inc. |
| | 1194 North Mathilda Ave. |
| | Sunnyvale, CA 94089-1206 |
| | US |
| Phone: | |
| Email: | gregory.ietf@gmail.com |
| | |
| | |
| Phone: | |
| Email: | |