

KARP Working Group  
Internet-Draft  
Intended status: Informational  
Expires: November 11, 2012

G. Lebovitz  
  
M. Bhatia  
Alcatel-Lucent  
May 10, 2012

**Keying and Authentication for Routing Protocols (KARP) Overview,  
Threats, and Requirements  
draft-ietf-karp-threats-reqs-05**

**Abstract**

Different routing protocols exist and each employs its own mechanism for securing the protocol packets on the wire. While most already have some method for accomplishing cryptographic message authentication, in many cases the existing methods are dated, vulnerable to attack, and employ cryptographic algorithms that have been deprecated. The "Keying and Authentication for Routing Protocols" (KARP) effort aims to overhaul and improve these mechanisms.

This document does not contain protocol specifications. Instead, it defines the areas where protocol specification work is needed and a set of requirements for KARP design teams to follow. [RFC 6518](#), "Keying and Authentication for Routing Protocols (KARP) Design Guidelines" is a companion to this document; KARP design teams will use them together to review and overhaul routing protocols. These two documents reflect the input of both the IETF's Security Area and Routing Area in order to form a mutually agreeable work plan.

This document has three main parts. The first part provides an overview of the KARP effort. The second part lists the threats from [RFC 4593](#), Generic Threats To Routing Protocols, that are in scope for attacks against routing protocols' transport systems, including any mechanisms built into the routing protocols themselves, which accomplish packet authentication. The third part enumerates the requirements that routing protocol specifications must meet when addressing those threats for [RFC 6518](#)'s "Work Phase 1", the update to a routing protocol's existing transport security.

**Status of this Memo**

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute

working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on November 11, 2012.

#### Copyright Notice

Copyright (c) 2012 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 Simplified BSD License.



## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">4</a>
<a href="#">1.1.</a>	Terminology . . . . .	<a href="#">5</a>
<a href="#">1.2.</a>	Requirements Language . . . . .	<a href="#">8</a>
<a href="#">2.</a>	KARP Effort Overview . . . . .	<a href="#">9</a>
<a href="#">2.1.</a>	KARP Scope . . . . .	<a href="#">9</a>
<a href="#">2.2.</a>	Incremental Approach . . . . .	<a href="#">10</a>
<a href="#">2.3.</a>	Goals . . . . .	<a href="#">11</a>
<a href="#">2.4.</a>	Non-Goals . . . . .	<a href="#">13</a>
<a href="#">2.5.</a>	Audience . . . . .	<a href="#">14</a>
<a href="#">3.</a>	Threats . . . . .	<a href="#">15</a>
<a href="#">3.1.</a>	Threat Sources . . . . .	<a href="#">15</a>
<a href="#">3.1.1.</a>	OUTSIDERS . . . . .	<a href="#">15</a>
<a href="#">3.1.2.</a>	Stolen Keys . . . . .	<a href="#">16</a>
<a href="#">3.1.2.1.</a>	Terminated Employee . . . . .	<a href="#">17</a>
<a href="#">3.2.</a>	Threat Actions In Scope . . . . .	<a href="#">18</a>
<a href="#">3.3.</a>	Threat Actions Out of Scope . . . . .	<a href="#">19</a>
<a href="#">4.</a>	Requirements for KARP Work Phase 1, the Update to a Routing Protocol's Existing Transport Security . . . . .	<a href="#">21</a>
<a href="#">5.</a>	Security Considerations . . . . .	<a href="#">27</a>
<a href="#">6.</a>	IANA Considerations . . . . .	<a href="#">28</a>
<a href="#">7.</a>	Acknowledgements . . . . .	<a href="#">29</a>
<a href="#">8.</a>	References . . . . .	<a href="#">30</a>
<a href="#">8.1.</a>	Normative References . . . . .	<a href="#">30</a>
<a href="#">8.2.</a>	Informative References . . . . .	<a href="#">30</a>
	Authors' Addresses . . . . .	<a href="#">32</a>



## **1. Introduction**

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 [[RFC4948](#)]. [Section 8.1](#) of that document states "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:

- o Create secure mechanisms and practices for operating routers.
- o Clean up the Internet Routing Registry repository (IRR), and securing both the database and the access, so that it can be used for routing verification.
- o Create specifications for cryptographic validation of routing message content.
- o Secure the routing protocols' packets on the wire

The first bullet is being addressed in the OPSEC working group. The second bullet should be addressed through liaisons with those running the IRR's globally. The third bullet is being addressed in the SIDR working group.

This document addresses the last item in the list above, securing the transmission of routing protocol packets on the wire, or rather securing the routing protocols' transport systems, including any mechanisms built into the routing protocols themselves which accomplish packet authentication. This effort is referred to as Keying and Authentication for Routing Protocols, or "KARP". KARP is concerned with issues and techniques for protecting the messages and their contents between directly communicating peers. This may overlap with, but is strongly distinct from, protection designed to ensure that routing information is properly authorized relative to sources of information. Such assurances are provided by other mechanisms and are outside the scope of this document and work that relies on it.

This document is one of two that together form the guidance and instructions for KARP design teams working to overhaul routing protocol transport security. The other document is the KARP Design Guide [[RFC6518](#)].

This document does not contain protocol specifications. Instead, its goal is to define the areas where protocol specification work is



needed and to provide a set of requirements for KARP design teams to follow as they tackle [\[RFC6518\]](#), [Section 4.1](#)'s "Work Phase 1", the update to a routing protocol's existing transport security.

This document has three main parts. The first part, found in [Section 2](#), provides an overview of the KARP effort. [Section 3](#) lists the threats from [\[RFC4593\]](#), Generic Threats To Routing Protocols, that are in scope for routing protocols' transport systems' per packet authentication. Therefore, this document does not contain a complete threat model; it simply points to the parts of the governing threat model that KARP design teams must address, and explicitly states which parts are out of scope for KARP design teams. [Section 4](#) enumerates the requirements that routing protocol specifications must meet when addressing those threats related to KARP's "Work Phase 1", the update to a routing protocol's existing transport security. ("Work Phase 2", a framework and usage of a KMP, will be addressed in a future document[s]).

This document uses the terminology "on the wire" to refer to the information used by routing protocols' transport systems. This term is widely used in IETF RFCs, but is used in several different ways. In this document, it is used to refer both to information exchanged between routing protocol instances, and to underlying protocols that may also need to be protected in specific circumstances. Individual protocol analysis documents will need to be more specific in their usage."

### [1.1](#). Terminology

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.

#### Identifier

The type and value used by a peer of an authenticated message exchange to signify who it is to another peer. The Identifier is used by the receiver as an 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 Authentication





Once the identity is decided, then there must be a cryptographic proof of that identity, that the peer really is who it asserts to be. Proof of identity can be arranged among peers in a few ways, for example symmetric and asymmetric pre-shared keys, or an asymmetric key contained in a certificate. Certificates can be used in ways that requires no additional supporting systems external to the routers themselves. An example of this would be using self signed certificates and a flat file list of "approved thumbprints". The use of these different identity authentication mechanisms 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.

#### KDF (Key derivation function)

A KDF is a function in which an input key and other input data is used to generate (or derive) keying material that can be employed by cryptographic algorithms. The key that is input to a KDF is called a key derivation key. KDFs can be used to generate one or more keys from either (i) a truly random or pseudorandom seed value or (ii) result of the Diffie-Hellman exchange or (iii) a non-uniform random source or (iv) a pre-shared key which may or may not be memorable by a human.

#### KMP (Key Management Protocol)

A protocol to establish a shared symmetric key between a pair (or a group) of users. It determines how secret keys are generated and made available to both the parties. If session or traffic keys are being used, KMP is responsible for generating them and determining when they should be renewed.

A KMP is helpful because it negotiates unique, random keys without administrator involvement. It also negotiates, as mentioned earlier, several of the SA parameters required for the secure connection, including key life times. It keeps track of those lifetimes, and negotiates new keys and parameters before they expire, again, without administrator interaction. Additionally, in the event of a security breach, changing KMP authentication credentials will immediately cause a rekey to occur for the Traffic Keys, and new Traffic Keys will be installed and used in the current connection.



## KMP Function

Any actual KMP used in the general KARP solution framework

## Peer Key

Keys that are used among peers as a basis for identifying one another. These keys may or may not be connection-specific, depending on how they were established, and what forms of identity and identity authentication mechanism used in the system. A peer key generally would be provided by a KMP that would later be used to derive fresh traffic keys.

## PRF

In cryptography, a pseudorandom function, abbreviated PRF, is a collection of efficiently-computable functions which emulate a random oracle in the following way: No efficient algorithm can distinguish (with significant advantage) between a function chosen randomly from the PRF family and a random oracle (a function whose outputs are determined at random). Informally, a PRF takes a secret key and a set of input values and produces random-seeming output values for each input value.

## PSK (Pre-Shared Key)

A key used to communicate with one or more peers in a secure configuration. Always distributed 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.

## SA (Security Association)

A relationship established between two or more entities to enable them to protect data they exchange. Examples of items that may exist in an SA include: Identifier, PSK, Traffic Key, cryptographic algorithms, key lifetimes.

## Traffic Key



The key (or one of a set of keys) used for protecting the routing protocol traffic. Since the traffic keys used in a particular connection are not a fixed part of a device configuration no 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.

## **1.2. Requirements Language**

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](#) [[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](#) [[RFC2119](#)].



## **2. KARP Effort Overview**

### **2.1. KARP Scope**

Three basic services may be employed in order to secure any piece of data as it is transmitted over the wire: confidentiality, authenticity, or integrity. The focus for the KARP working group will be message authentication and message integrity only. This work explicitly excludes, at this point in time, privacy services. Non-repudiation is also excluded as a goal at this time. Since the objective of most routing protocols is to broadly advertise the routing topology, routing protocol packets are commonly sent in the clear; confidentiality is not normally required for routing protocols. However, ensuring that routing peers are authentically identified, and that no rogue peers or unauthenticated packets can compromise the stability of the routing environment is critical, and thus our focus. Confidentiality and non-repudiation may be addressed in future work.

OSPF [[RFC5709](#)], IS-IS [[RFC5310](#)], LDP [[RFC5036](#)], and RIP [[RFC2453](#)] [[RFC4822](#)] already have existing mechanisms for cryptographically authenticating and integrity checking the messages on the wire. Products with these mechanisms have been produced, code has been written, and both have been optimized for these existing security mechanisms. Rather than turn away from these mechanisms, this document aims to enhance them, updating them to modern and secure levels.

Therefore, the scope of KARP's roadmap of work includes:

- o Making use of existing routing protocol transport security mechanisms, where they exist, and enhancing or updating them as necessary for modern cryptographic best practices. [[RFC6518](#)], [Section 4.1](#) labels this KARP's "Work Phase 1."
- 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. [[RFC6518](#)], [Section 4.1](#) labels this KARP's "Work Phase 2."
- o Specifying an automated key management protocol that may be combined with the bits-on-the-wire mechanisms. [[RFC6518](#)], [Section 4.1](#) labels this KARP's "Work Phase 2."

Neither this document nor [[RFC6518](#)] contain protocol specifications. Instead, they define the areas where protocol specification work is needed and set a direction, a set of requirements, and priorities for addressing that specification work.





There are a set of threats to routing protocols that are considered in-scope for KARP, and a set considered out-of-scope. These are described in detail in the Threats ([Section 3](#)) section below.

## **2.2. Incremental Approach**

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. The principle of "crawl, walk, run" will be employed. Thus routing protocol authentication mechanisms may not go immediately from their current state to a state reflecting the best possible, most modern security practices. This point is important as 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.

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 a future release of the specification. This approach 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 were at play in an incremental improvement roadmap was seen in the improvement of BGP's [\[RFC4271\]](#) security via the TCP Authentication Option (TCP-AO) [\[RFC5925\]](#) effort. It would have been ideal, and reflected best common security practice, to have a fully specified key management protocol for negotiating TCP-AO's keying material, e.g., using certificates for



peer authentication. However, in the spirit of incremental deployment, we first addressed issues like cryptographic algorithm agility, replay attacks, and TCP session resetting in the base TCP-AO protocol, and then later began work to layer key management on top of it.

### **2.3. Goals**

The goals and general guidance for the KARP work follow.

1. Provide authentication and integrity protection for messages on the wire of existing routing protocols.
2. Define a path to incrementally improve security of the routing infrastructure as explained in the earlier sections.
3. Ensure that the improved security solutions on currently running routing infrastructure equipment are deployable. This begs the consideration of the current state of processing power available on routers in the network today.
4. Operational deployability - A solution's acceptability will also be measured by how deployable the solution is by common operator teams using common deployment processes and infrastructures. Specifically, we will try to make these solutions fit as well as possible into current operational practices and 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 and integrity mechanisms in their networks, 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 [[RFC3562](#)] on their exterior Border Gateway Protocol (eBGP) [[ISR2008](#)] sessions. Over 55% also deploy TCP-MD5 on their interior Border Gateway Protocol (iBGP) connections, and 50% make use of TCP-MD5 offered on some other internal gateway protocol (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 TCP-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 [[RFC4301](#)] to protect the routing protocol, and this was a decline from the few that reported doing so in the previous year's report. From our personal conversations with operators, of those using MD5, almost all report using one, manually-distributed key throughout the entire network. These same operators report that the 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 protecting BGP sessions 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, the 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. Key change is perceived as a relatively low priority compared to other operational issues.
- E. Lack of staff to implement the changes on a device-by-device basis.
- F. There are three use cases for operational peering at play here: peers and interconnection with other operators, iBGP, and other routing sessions within a single operator, and operator-to-customer 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 (CPE). The same operator reported that if the customer mandated it, 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 several) keys used for authentication in these environments could easily wage an attack. Alternatively, the operator could offer the keys to others who would wage the attack. In either case, the attacker could then bring down many of the adjacencies, causing destabilization to the routing system.

5. Whatever mechanisms KARP specifies need to be easier to deploy than the current methods, 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.
6. Address the threats enumerated below in the "Threats" section ([Section 3](#)) for each routing protocol. 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.
7. Create a re-usable architecture, framework, and guidelines for various IETF working groups who will address these security improvements for various Routing Protocols. The crux of the KARP work is to re-use the architecture, guidelines and the 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.
8. Bridge any gaps between IETF's Routing and Security Areas by recording agreements on work items, roadmaps, and guidance from the cognizant Area Directors and the Internet Architecture Board (IAB).

#### **[2.4.](#) Non-Goals**

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

- o Confidentiality of the packets on the wire. Once this roadmap is realized, we may revisit work on privacy.
- o Message content validity (routing database validity). This work is being addressed in other IETF efforts, like SIDR.





## **2.5. Audience**

The audience for this document 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 the 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 newly specified 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 only if 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. As a result, the GROW and OPSEC WGs should be kept squarely in the loop as well.



### **3. Threats**

In this document we will use the definition of "threat" as defined in [RFC4949](#) [[RFC4949](#)]: "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 the KARP effort. It also lists those threats that are explicitly out of scope for the KARP effort.

This document leverages the "Generic Threats to Routing Protocols" model, [[RFC4593](#)]. Specifically, the threats below were derived by reviewing [[RFC4593](#)], analyzing the KARP problem space relative to it, and simply listing the threats that are applicable to the KARP design teams' work. This document categorizes [[RFC4593](#)] threats into those in scope and those out of scope for KARP. Each in-scope threat is discussed below, and its applicability to the KARP problem space is described. As such, the below text intentionally does not constitute a self-standing, complete threat analysis, but rather describes the applicability of the existing threat analysis [[RFC4593](#)] relevant to KARP.

Note: terms from [[RFC4593](#)] appear capitalized below -- e.g. OUTSIDERS -- so as to make explicit the term's origin, and to enable rapid cross referencing to the source RFC.

For convenience, a terse definition of most [[RFC4593](#)] terms is offered here. Those interested in a more thorough description of routing protocol threat sources, motivations, consequences and actions will want to read [[RFC4593](#)] before continuing here.

#### **3.1. Threat Sources**

##### **3.1.1. OUTSIDERS**

One of the threats that will be addressed in this roadmap are those where the source is an OUTSIDER. An OUTSIDER attacker 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. OUTSIDERS are not legitimate participants in the routing protocol. The use of message authentication and integrity protection specifically aims to identify packets originating from OUTSIDERS.

KARP design teams will consider two specific use cases of OUTSIDERS: those on-path, and those off-path.



- o On-Path - These sources have control of a network resource or a tap that sits along the path of packets between the two routing peers. A "Man-in-the-Middle" (MitM) is an on-path attacker. From this vantage point, the attacker can conduct either active or passive attacks. An active attack occurs when the attacker actually places packets on the network as part of the attack. One active MitM attack relevant to KARP, an active wiretapping attack, occurs when the attacker tampers with packets moving between two legitimate router peers 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 detect MitM attacks and reject packets from such attacks (i.e. treat them as not authentic). Passive on-path attacks occur when the attacker silently gathers data and analyses it to gain advantage. Passive activity by an on-path attacker may often eventually lead to an active attack.
- o Off-Path - These sources sit on some network outside of that over which runs the packets between two routing peers. The source may be one or several hops away. Off-path attackers can launch active attacks, such as SPOOFING or denial-of-service (DoS) attacks, to name a few.

### **3.1.2. Stolen Keys**

This threat source exists when an unauthorized entity somehow manages to gain access to keying material. Using this material, the attacker could send packets that pass the authenticity checks based on message authentication codes (MACs). The resulting traffic might appear to come from router A to router B, and thus the attacker could impersonate an authorized peer. The attacker could then adversely affect network behavior by sending bogus messages that appear to be authentic. The attack source possessing the stolen keys could be on-path, off-path, or both.

The obvious mitigation for stolen keys is to change the keys currently in use by the legitimate routing peers. This mitigation can be either reactive or pro-active. Reactive mitigation occurs when keys are changed only after having discovered that the previous keys fell into the possession of unauthorized users. The stolen keys, reactive mitigation case is highlighted here in order to explain a common operational situation where new keying material will become necessary with little or no advanced warning. In such a case new keys must be able to be installed and put into use very quickly, and with little operational expense. Pro-active mitigation occurs when an operator assumes that unauthorized possession will occur from time to time without being discovered, and the operator moves to new



keying material in order to cut short, or make nonexistent, an attacker's window of opportunity to use the stolen keys effectively.

In KARP, we can address the attack source with stolen keys by creating specifications that make it practical for the operator to quickly change keys without disruption to the routing system, and with minimal operational overhead. Operators can further mitigate the stolen keys case by habitually changing keys.

#### **3.1.2.1. Terminated Employee**

A terminated employee is an important example of a "stolen keys" threat source to consider. Staff attrition is a reality in routing operations, and so regularly causes the potential for a threat source. The threat source risk arises when a network operator who had been granted access to keys ceases to be an employee. If new keys are deployed immediately, the situation of a terminated employee can become a "stolen keys, pro-active" case, as described above, rather than a "stolen keys, reactive" case.

On one hand, terminated employees could be considered INSIDERS rather than OUTSIDERS, because at one point in time they were authorized to have the keys. On the other hand, they aren't really a BYZANTINE attacker, which is defined to be an attack from an INSIDER, a legitimate router. Further, once terminated, the authorization granted to the terminated employee regarding the keys is revoked. If they maintain possession of the keys they are acting in an unauthorized way. If they go on to use those keys to launch an attack they are definitely acting in an unauthorized way. In this way the terminated employee becomes an OUTSIDER at the point of termination, they cease to be legitimate participants in the routing system. It behooves the operator to change the keys, to enforce the revocation of authorization of the old keys, in order to minimize the threat source's window of opportunity.

Regardless of whether one considers a terminated employee an "insider" or an OUTSIDER, it is important to consider them a threat source, study the use case, and address the threats therein. In such a case within the KARP context, new keys must be able to be installed and made operational in the routing protocols very quickly, with zero impact to the routing system, and with little operational expense.

The threat source of the terminated employee and/or the detected-stolen-keys drives the requirement for quick and easy key rollover. The threat actions associated with these sources are mitigated if the operator has mechanisms in place (both inherent in the protocol, as well as built into their management systems) that allow them to roll the keys quickly with minimal impact to the routing system, at low





operational cost.

### **3.2. Threat Actions In Scope**

These ATTACK ACTIONS are in scope for KARP:

- 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 DoS attacks at the transport layer - This is an example of SPOOFING. It can also be an example of FALSIFICATION and INTERFERENCE (see below). It occurs when an attacker sends spoofed packets aimed at halting or preventing the underlying protocol over which the routing protocol runs. For example, BGP running over TLS will still not solve the problem of being able to send a spoofed TCP FIN or TCP RST and causing the BGP session to go down. Since this attack depends on spoofing, operators are encouraged to deploy proper authentication mechanisms to prevent such attacks. Specification work should ensure that Routing Protocols can operate over transport sub-systems in a fashion that is resilient to such DoS attacks.
- 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, or sending entirely false messages. FALSIFICATION from BYZANTINES (see the Threats Out of Scope section 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:
  - A. ADDING NOISE
  - B. REPLAYING OUT-DATED PACKETS
  - C. INSERTING MESSAGES



## D. CORRUPTING MESSAGES

## E. BREAKING SYNCHRONIZATION

## F. Changing message content

- o DoS attacks using the authentication mechanism - This includes an attacker sending packets that confuse or overwhelm a security mechanism itself. An example is initiating an overwhelming load of spoofed routing protocol packets that contain a MAC, so that the receiver needs to spend the processing cycles to check the MAC, only to discard the spoofed packet, consuming substantial CPU resources. Another example is when an attacker sends an overwhelming load of keying protocol initiations from bogus sources.
- 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. Using per session keys is another widely used method for reducing the number of brute force attacks as this would make it difficult to guess the keys.

### **3.3. Threat Actions Out of Scope**

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 ([Section 2.1](#)) that may be levied by a BYZANTINE source are therefore also out of scope, e.g. FALSIFICATION, or unauthorized message content by a legitimate authorized peer.

In addition, these other attack actions are out of scope for this work:

- o SNIFFING - passive observation of route message contents in flight. Data privacy, as achieved by data encryption, is the common mechanism for preventing SNIFFING. While useful, especially to prevent the gathering of data needed to perform an off-path packet injection attack, data encryption is out-of-scope for KARP.



- o INTERFERENCE due to:
  - A. NOT FORWARDING PACKETS - cannot be prevented with cryptographic authentication. Note: If sequence numbers with sliding windows are used in the solution (as is done, for example, in IPsec's ESP [[RFC4303](#)] and BFD [[RFC5880](#)], a receiver can at least detect the occurrence of this attack.
  - B. DELAYING MESSAGES - cannot be prevented with cryptographic authentication. Note: Timestamps can be used to detect delays.
  - C. DENIAL OF RECEIPT - cannot be prevented with cryptographic authentication
  - D. UNAUTHORIZED MESSAGE CONTENT - the work of the IETF's SIDR working group (<http://www.ietf.org/html.charters/sidr-charter.html>).
  - E. DoS attacks not involving the routing protocol. 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 KARP will address. Many such examples could be contrived.



#### **4. Requirements for KARP Work Phase 1, the Update to a Routing Protocol's Existing Transport Security**

The KARP Design Guide [\[RFC6518\]](#), [Section 4.1](#) describes two distinct work phases for the KARP effort. This section addresses requirements for the first work phase only, "Work Phase 1", the update to a routing protocol's existing transport security. "Work Phase 2", a framework and usage of a KMP, will be addressed in a future document(s)."

The following list of requirements SHOULD be addressed by a KARP Work Phase 1 security update to any Routing Protocol (according to [section 4.1](#) of the KARP Design Guide [\[RFC6518\]](#) document). IT IS RECOMMENDED that any Work Phase 1 security update to a Routing Protocol contain a section of the specification document that describes how each of the below requirements are met. It is further RECOMMENDED that justification be presented for any requirements that are NOT addressed.

1. Clear definitions of which elements of the transmitted data (frame, packet, segment, etc.) are protected by the authentication mechanism
2. Strong cryptographic algorithms, as defined and accepted by the IETF security community, MUST be specified. 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 because research identifying weaknesses in cryptographic algorithms can cause the security community to reduce confidence in some algorithms. Breaking a cipher isn't a matter of if, but when it will occur. Having the ability to specify alternate algorithms (algorithm agility) within the protocol specification to support such an event is essential. Mandating two algorithms provides both a redundancy, and a mechanism for enacting that redundancy when needed. Further, the mechanism MUST describe the generic interface for new cryptographic algorithms to be used, so that implementers can use algorithms other than those specified, and so that new algorithms may be specified and supported in the future.
4. Secure use of PSKs, offering both operational convenience and a baseline level of security, MUST be specified.





5. Routing protocols (or the transport or network mechanism protecting routing protocols) should be able to detect and reject replayed messages. For non TCP based protocols like OSPF [RFC2328], IS-IS [RFC1195], etc., two routers are said to have a session up if they are able to exchange protocol packets. Packets captured from one session must not be able to be re-sent and accepted during a later session. Additionally, replay mechanisms must work correctly even in the presence of routing protocol packet prioritization by the router.
  - A. There is a specific case of replay attack combined with spoofing that must be addressed. In several routing protocols (e.g., OSPF [RFC2328], IS-IS [RFC1195], BFD [RFC5880], RIP [RFC2453], etc.), all speakers share the same key (K) on a broadcast segment. The ability to run a MAC operation with K is used for identity validation, and (currently) no other identity validation check is performed. Assume there are four routers using authentication on a LAN, R1 - R4. Also assume attacker "Z", who is NOT a legitimate neighbor, is observing and recording packets on the same LAN segment. Z captures a packet from R1, and changes the source IP, spoofing it to that of R2, then sends the packet on the LAN. Z does not have K, but in this case it does not matter because R1 already performed the MAC operation, and Z simply re-uses that MAC. R3 and R4 will process the packet as if coming from R2, the MAC check will return valid, and they will update their route tables accordingly. R3 and R4 have confirmed that the MAC was created by someone holding K, but not that it was actually sent by R2. This is a well known attack with known solutions. Some string must be added into the MAC operation that uniquely identifies the sender. Said string must also be located in the packet such that if that string were to be altered after the MAC operation, it would be detected by the receiver. Examples of solutions used in other protocols include sequence numbers with sliding acceptance windows, time stamps, IP header info (SRC, DST), unique identifiers which are temporarily bound to an IP Address.
6. A change of security parameters REQUIRES, and even forces, a change of session traffic keys. The specific security parameters for the various routing protocols will differ, and will be defined by each protocols design team. Some examples may include: master key, key lifetime, cryptographic algorithm, etc. If one of these configured parameters changes, then a new session traffic key must immediately be established using the updated parameters. The routing protocol security mechanisms MUST support this behavior.



7. Intra-session re-keying which occurs without a break or interruption to the current routing session, and, if possible, without data loss, MUST be specified. Keys need to be changed periodically, for operational confidentiality (e.g. when an administrator who had access to the keys leaves an organization) and for entropy purposes, and a re-keying mechanism enables the operators to execute the change without productivity loss.
8. Efficient re-keying SHOULD be provided. The specification SHOULD support rekeying during a session without needing to try/compute multiple keys on a given packet. The rare exception will occur if a routing protocols design team can find no other way to re-key and still adhere to the other requirements in this section.
9. New mechanisms must resist DoS attacks described as in-scope in [Section 3.2](#). Routers protect the control plane by implementing mechanisms to filter completely or rate limit traffic not required at the control plane level (i.e., unwanted traffic). Typically line rate packet filtering capabilities look at information at or below the IP and transport (TCP or UDP) headers, but do not include higher layer information. Therefore the new mechanisms shouldn't hide nor encrypt the information carried in the IP and transport layers in control plane packets.
10. Mandatory cryptographic algorithms and mechanisms MUST be specified for a routing protocol. Further, the protocol specification MUST define default security mechanism settings for all implementations to use when no explicit configuration is provided. To understand the need for this requirement, consider the case where a routing protocol mandates 3 different cryptographic algorithms for a MAC operation. If company A implements algorithm 1 as the default for this protocol, while company B implements algorithm 2 as the default, then two operators who enable the security mechanism with no explicit configuration other than a PSK will experience a connection failure. It is not enough that each implementation implement the 3 mandatory algorithms; one default must further be specified in order to gain maximum out-of-the-box interoperability.
11. For backward compatibility reasons manual keying MUST be supported.
12. Architecture of the specification SHOULD consider and allow for future use of a KMP.



13. The authentication mechanism in the Routing Protocol MUST be decoupled from the key management system used. It MUST be obvious how the keying material was obtained, and the process for obtaining the keying material MUST exist outside of the Routing Protocol. This will allow for the various key generation methods, like manual keys and KMPs, to be used with the same Routing Protocol mechanism.
14. Convergence times of the Routing Protocols SHOULD NOT be materially affected. "Materially" is defined here as anything greater than a 5% increase in convergence time. Changes in the convergence time will be immediately verifiable by convergence performance test beds already in use by most router vendors and service providers. 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 processing power increases by Moore's law over time, meaning that the same route calculations and table population routines will decrease in duration over time. Therefore, this requirement should be considered only in terms of total number of protocol packets that must be exchanged, and less for the computational intensity of processing any one message. Alternatively this can be simplified by saying that the new mechanisms should only result in a minimal increase in the number of routing protocol packets passed between the peers.
15. The changes to or addition of security mechanisms SHOULD NOT cause a refresh of route advertisements or cause additional route advertisements to be generated.
16. Router implementations provide prioritized treatment for certain protocol packets. For example, OSPF HELLO packets and ACKs are prioritized for processing above other OSPF packets. The security mechanism SHOULD NOT interfere with the ability to observe and enforce such prioritization. Any effect on such priority mechanisms MUST be explicitly documented and justified. Replay protection mechanisms provided by the routing protocols MUST work even if certain protocol packets are offered prioritized treatment.
17. Routing protocols MUST only send minimal information regarding the authentication mechanisms and the parameters in its protocol packets. One reason for this is to keep the Routing Protocols as clean and focused as possible, and load security negotiations into the future KMP as much as possible. Another reason is to avoid exposing any security negotiation information unnecessarily to possible attackers on the path.



18. Routing protocols that rely on the IP header (or information separate from routing protocol payload) to identify the neighbor that originated the packet, MUST either protect the IP header or provide some other means to authenticate the neighbor.  
[RFC6039] describes some attacks that are based on this.
19. Every new KARP-developed security mechanisms MUST support incremental deployment. It will not be feasible to deploy a new Routing Protocol authentication mechanism throughout a network instantaneously. It also may not be possible to deploy such a mechanism to all routers in a large autonomous system (AS) at one time. Proposed solutions MUST support an incremental deployment method that provides some benefit for those who participate. Because of this, there are several requirements that any proposed KARP mechanism should consider.
  - A. The Routing Protocol security mechanism MUST enable each router to configure use of the security mechanism on a per-peer basis where the communication is peer-to-peer (unicast).
  - B. Every new KARP-developed security mechanism MUST provide backward compatibility in the message formatting, transmission, and processing of routing information carried through a mixed security environment. Message formatting in a fully secured environment MAY be handled in a non-backward compatible fashion though care must be taken to ensure that routing protocol packets can traverse intermediate routers that don't support the new format.
  - C. In an environment where both secured and non-secured systems are interoperating, a mechanism MUST exist for secured systems to identify whether a peer intended the messages to be secured.
  - D. In an environment where secured service is in the process of being deployed, a mechanism MUST exist to support a transition free of service interruption (caused by the deployment per se).
20. The introduction of mechanisms to improve routing security may increase the processing performed by a router. Since most of the currently deployed routers do not have hardware to accelerate cryptographic operations, these operations could impose a significant processing burden under some circumstances. Thus proposed solutions should be evaluated carefully with regard to the processing burden they may impose, since deployment may be impeded if network operators perceive that a





solution will impose a processing burden which either incurs substantial capital expense, or threatens to destabilize routers.

21. Given the number of routers that would require the new authentication mechanisms in a typical ISP deployment, solutions can increase their appeal by minimizing the burden imposed on all routers in favor of confining significant work loads to a relatively small number of devices. Optional features or increased assurance that engenders more pervasive processing loads MAY be made available for deployments where the additional resources are economically justifiable.
22. New authentication and security mechanisms should not rely on systems external to the routing system (the equipment that is performing forwarding) in order for the routing system to function. In order to ensure the rapid initialization and/or return to service of failed nodes it is important to reduce reliance on these external systems to the greatest extent possible. Proposed solutions SHOULD NOT require connections to external systems, beyond those directly involved in peering relationships, in order to return to full service. It is however acceptable for the proposed solutions to require post initialization synchronization with external systems in order to fully synchronize the security information.

If authentication and security mechanisms rely on systems external to the routing system, then there MUST be one or more options available to avoid circular dependencies. It is not acceptable to have a routing protocol (e.g., unicast routing) depend upon correct operation of a security protocol that, in turn, depends upon correct operation of the same instance of that routing protocol (i.e., the unicast routing). However, it is okay to have operation of a routing protocol (e.g., multicast routing) depend upon operation of a security protocol, which depends upon an independent routing protocol (e.g., unicast routing). Similarly it would be okay to have the operation of a routing protocol depend upon a security protocol, which in turn uses an out of band network to exchange information with remote systems.



## 5. Security Considerations

This document is mostly about security considerations for the KARP efforts, both threats and requirements for addressing those threats. More detailed security considerations were placed in the Security Considerations section of the KARP Design Guide [[RFC6518](#)]document.

Spoofing by a Legitimate Neighbor - In several routing protocols (e.g OSPF) all speakers share the same key, a group key, on a broadcast segment. Possession of the group key itself is used for identity validation, and no other identity check is used. Under these conditions an attack exists where one neighbor (E.g. Router 1, or R1) can masquerade as a different neighbor, R2, by sending spoofed packets using R2 as the source IP address. When other neighbors, R3 and R4, receive these packets, they will calculate the MAC successfully, and process its contents as if it originated from R2. SPOOFING this way, the attacker can succeed in several different types of attacks, including FALSIFICATION and INTERFERENCE. The source of such an attack is a BYZANTINE actor, since the attack originates from a legitimate actor in the routing system, and such sources are out of scope for KARP. This type of attack has been well documented in the group keying problem space, and it's non-trivial to solve. The common method used to prevent this type of attack is to use a unique key for each sender rather than a group key. Other solutions exist within the group keying realm, but they come with significant increases in complexity and computational intensity.



## **6. IANA Considerations**

This document has no actions for IANA.

## **7. Acknowledgements**

The majority of the text for version -00 of this document was taken from "Roadmap for Cryptographic Authentication of Routing Protocol Packets on the Wire", [draft-lebovitz-karp-roadmap](#), authored by Gregory M. Lebovitz.

Brian Weis provided significant assistance in handling the many comments that came back during IESG review.

We would like to thank the following people for their thorough reviews and comments: Brian Weis, Yoshifumi Nishida, Stephen Kent, Vishwas Manral.

Author Gregory M. Lebovitz was employed at Juniper Networks, Inc. for the majority of the time he worked on this document, though not at the time of its publishing. Thus Juniper sponsored much of this effort.





## **8. References**

### **8.1. Normative References**

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4593] Barbir, A., Murphy, S., and Y. Yang, "Generic Threats to Routing Protocols", [RFC 4593](#), October 2006.
- [RFC4948] Andersson, L., Davies, E., and L. Zhang, "Report from the IAB workshop on Unwanted Traffic March 9-10, 2006", [RFC 4948](#), August 2007.

### **8.2. Informative References**

- [ISR2008] McPherson, D. and C. Labovitz, "Worldwide Infrastructure Security Report", October 2008, <[http://www.arbornetworks.com/dmdocuments/ISR2008\\_US.pdf](http://www.arbornetworks.com/dmdocuments/ISR2008_US.pdf)>.
- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", [RFC 1195](#), December 1990.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), April 1998.
- [RFC2453] Malkin, G., "RIP Version 2", STD 56, [RFC 2453](#), November 1998.
- [RFC3562] Leech, M., "Key Management Considerations for the TCP MD5 Signature Option", [RFC 3562](#), July 2003.
- [RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", [RFC 4301](#), December 2005.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", [RFC 4303](#), December 2005.
- [RFC4822] Atkinson, R. and M. Fanto, "RIPv2 Cryptographic Authentication", [RFC 4822](#), February 2007.
- [RFC4949] Shirey, R., "Internet Security Glossary, Version 2", [RFC 4949](#), August 2007.
- [RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", [RFC 5036](#), October 2007.



- [RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", [RFC 5310](#), February 2009.
- [RFC5709] Bhatia, M., Manral, V., Fanto, M., White, R., Barnes, M., Li, T., and R. Atkinson, "OSPFv2 HMAC-SHA Cryptographic Authentication", [RFC 5709](#), October 2009.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", [RFC 5880](#), June 2010.
- [RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", [RFC 5925](#), June 2010.
- [RFC6039] Manral, V., Bhatia, M., Jaeggli, J., and R. White, "Issues with Existing Cryptographic Protection Methods for Routing Protocols", [RFC 6039](#), October 2010.
- [RFC6518] Lebovitz, G. and M. Bhatia, "Keying and Authentication for Routing Protocols (KARP) Design Guidelines", [RFC 6518](#), February 2012.



Authors' Addresses

Gregory Lebovitz  
Aptos, California 95003  
USA

Email: [gregory.ietf@gmail.com](mailto:gregory.ietf@gmail.com)

Manav Bhatia  
Alcatel-Lucent  
Bangalore,  
India

Phone:

Email: [manav.bhatia@alcatel-lucent.com](mailto:manav.bhatia@alcatel-lucent.com)

