

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 17, 2014

R. Schatzmayr
Deutsche Telekom AG
G. Heron, Ed.
M. Konstantynowicz, Ed.
M. Townsley
Cisco Systems
July 16, 2013

Keyed IPv6 Tunnel
draft-mkonstan-keyed-ipv6-tunnel-00

Abstract

This document describes a simple L2 Ethernet over IPv6 tunnel encapsulation with mandatory 64-bit authentication key for connecting L2 Ethernet attachment circuits identified by IPv6 addresses. The encapsulation is based on L2TPv3 over IP.

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].

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 January 17, 2014.

Copyright Notice

Copyright (c) 2013 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

1. Introduction	2
2. Static 1:1 Mapping Without a Control Plane	2
3. 64-bit Authentication Key	3
4. Encapsulation	3
5. IANA Considerations	6
6. Security Considerations	6
7. Acknowledgements	7
8. References	7
8.1. Normative References	7
8.2. Informative References	7
Authors' Addresses	7

1. Introduction

L2TPv3, as defined in RFC3931 [RFC3931], provides a dynamic mechanism for tunneling Layer 2 (L2) "circuits" across a packet-oriented data network (e.g., over IP), with multiple attachment circuits multiplexed over a single pair of IP address endpoints (i.e. a tunnel) using the L2TPv3 session ID as a circuit discriminator.

Implementing L2TPv3 over IPv6 provides the opportunity to utilize unique IPv6 addresses to identify Ethernet attachment circuits directly, leveraging the key property that IPv6 offers, a vast number of unique IP addresses. In this case, processing of the L2TPv3 Session ID may be bypassed upon receipt as each tunnel has one and only one associated session. This local optimization does not hinder the ability to continue supporting the multiplexing of circuits via the Session ID on the same router for other L2TPv3 tunnels.

2. Static 1:1 Mapping Without a Control Plane

Use of the L2TPv3 Control Plane is optional. When the control plane is not used, local configuration creates a one-to-one mapping between the access-side L2 attachment circuit and the IP address used in the network-side IPv6 encapsulation. Further, circuit monitoring is performed using Ethernet OAM mechanisms (802.1ag and/or Y.1731).

The L2TPv3 encapsulating router identifies each Ethernet L2 attachment circuit by the Ethernet VLAN stack present on Ethernet frames on the access side

- o port mode access - physical port identifies a L2 attachment circuit.
- o single-stack access - s-tag or c-tag with S-VID or C-VID value identifies a L2 attachment circuit.
- o multi-stack access - (s-tag, c-tag) with tuple (S-VID, C-VID) identifies a L2 attachment circuit.

L2 attachment connection identifiers s-tag or (s-tag, c-tag) are treated with local significance and are not required to be forwarded over the IPv6 network (though the operator may prefer to forward tags in some cases).

The L2TPv3 encapsulating router identifies each L2TPv3 tunnel endpoint by a distinct /128 IPv6 address in the packet header of L2TPv3 IPv6 packets received and transmitted on the network side.

In the event that an IPv6 address used in L2TPv3 does not directly correspond to one and only one attachment circuit on both sides of the L2TPv3 tunnel, the Session ID may be used for additional granularity. This allows for other addressing schemes that may require additional bits beyond those which can fit in the IPv6 header address field.

3. 64-bit Authentication Key

All packets MUST carry a 64-bit authentication key in the L2TPv3 cookie field. The cookie MUST be 64-bits long in order to provide sufficient protection against a brute force blind insertion attack.

In absence of the L2TPv3 Control Plane, the L2TPv3 encapsulating router must be provided with local configuration of the 64-bit authentication cookie for each local and remote IPv6 endpoint - note that cookies are asymmetric, so local and remote endpoints may send different cookie values. The value of the cookie must be able to be changed at any time in a manner that does not drop any legitimate tunneled packets - i.e. the receiver must be willing to accept both "old" and "new" cookie values during a change of cookie value.


4. Encapsulation

RFC4719 [RFC4719] describes encapsulation of Ethernet over L2TPv3. Paraphrasing from this document, the Ethernet frame, without the

preamble or frame check sequence (FCS), is encapsulated in L2TPv3 and is sent as a single packet by the ingress router.

The s-tag (or in the multi-stack access case the s-tag and c-tag) SHOULD be removed before the packet is encapsulated.

The full encapsulation is as follows:



The combined IPv6 and L2TPv3 header contains the following fields:

- o Ver. Set to 0x6 to indicate IPv6.
- o Traffic Class. May be set by the ingress router to ensure correct PHB treatment by transit routers between the ingress and egress, and correct QoS disposition at the egress router.
- o Flow Label. May be set by the ingress router to indicate a flow of packets from the client which may not be reordered by the network (if there is a requirement for finer grained ECMP load balancing than per-circuit load balancing).
- o Payload Length. Set to the length of the packet, excluding the IPv6 header (i.e. the length from the Session ID to the end of the packet).
- o Next Header. Set to 0x73 to indicate that the next header is L2TPv3.
- o Hop Limit. Set to 0xFF, and decremented by one by each router in the path to the egress router.
- o Source Address. IPv6 source address for the tunnel. In the "Static 1:1" case the IPv6 source address may correspond to a port or VLAN being transported as an L2 circuit, or may be a loopback address terminating inside the router (e.g. if L2 circuits are being used within a multipoint VPN) or may be an anycast address terminating on a data center virtual machine.
- o Destination Address. IPv6 destination address for the tunnel. As with the source address this may correspond to a port or VLAN being transported as an L2 circuit or may be a loopback or anycast address.
- o Session ID. In the "Static 1:1 mapping" case described in Section 2, the IPv6 address resolves to an L2TPv3 session immediately, thus the Session ID may be ignored upon receipt. For compatibility with other tunnel termination platforms supporting only 2-stage resolution (IPv6 Address + Session ID), this specification recommends supporting explicit configuration of Session ID to any value other than zero. For cases where both tunnel endpoints support one-stage resolution (IPv6 Address only), this specification recommends setting the Session ID to all ones for easy identification in case of troubleshooting.
- o Cookie. 64 bits, configured and described as in Section 3. All packets for a destined L2 Circuit (or L2TPv3 Session) must match the configured Cookie value or be discarded (see RFC3931 [RFC3931] for more details).

- o Payload. The customer data, with s-tag or s-tag/c-tag removed. As noted above preamble and FCS are stripped before encapsulation. A new FCS will be added at each hop when the IP packet is transmitted.

5. IANA Considerations

None.

6. Security Considerations

Packet spoofing for any type of Virtual Private Network (VPN) tunneling protocol is of particular concern as insertion of carefully constructed rogue packets into the VPN transit network could result in a violation of VPN traffic separation, leaking data into a customer VPN. This is complicated by the fact that it may be particularly difficult for the operator of the VPN to even be aware that it has become a point of transit into or between customer VPNs.

Keyed IPv6 encapsulation provides traffic separation for its VPNs via use of separate 128-bit IPv6 addresses to identify the endpoints. The mandatory authentication key carried in the L2TPv3 cookie field, provides an additional check to ensure that an arriving packet is intended for the identified tunnel.

In the presence of a blind packet spoofing attack, the authentication key provides security against inadvertent leaking of frames into a customer VPN, like in case of L2TPv3 RFC3931 [RFC3931]. To illustrate the type of security that it is provided in this case, consider comparing the validation of a 64-bit Cookie in the L2TPv3 header to the admission of packets that match a given source and destination IP address pair. Both the source and destination IP address pair validation and Cookie validation consist of a fast check on cleartext header information on all arriving packets. However, since L2TPv3 uses its own value, it removes the requirement for one to maintain a list of (potentially several) permitted or denied IP addresses, and moreover, to guard knowledge of the permitted IP addresses from hackers who may obtain and spoof them. Further, it is far easier to change a compromised L2TPv3 Cookie than a compromised IP address," and a cryptographically random RFC4086 [RFC4086] value is far less likely to be discovered by brute-force attacks compared to an IP address.

For protection against brute-force, blind, insertion attacks, a 64-bit Cookie MUST be used with all tunnels.

Note that the Cookie provides no protection against a sophisticated man-in-the-middle attacker who can sniff and correlate captured data between nodes for use in a coordinated attack.

The L2TPv3 64-bit cookie must not be regarded as a substitute for security such as that provided by IPsec when operating over an open or untrusted network where packets may be sniffed, decoded, and correlated for use in a coordinated attack.

7. Acknowledgements

...

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.
- [RFC3931] Lau, J., Townsley, M., and I. Goyret, "Layer Two Tunneling Protocol - Version 3 (L2TPv3)", RFC 3931, March 2005.
- [RFC4086] Eastlake, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, June 2005.
- [RFC4719] Aggarwal, R., Townsley, M., and M. Dos Santos, "Transport of Ethernet Frames over Layer 2 Tunneling Protocol Version 3 (L2TPv3)", RFC 4719, November 2006.

8.2. Informative References

- [RFC1700] Reynolds, J. and J. Postel, "Assigned Numbers", RFC 1700, October 1994.

Authors' Addresses

Rainer Schatzmayr
Deutsche Telekom AG

Email: rainer.schatzmayr@telekom.de

Giles Heron (editor)
Cisco Systems

Email: giheron@cisco.com

Maciek Konstantynowicz (editor)
Cisco Systems

Email: maciek@cisco.com

Mark Townsley
Cisco Systems

Email: townsley@cisco.com