

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: May 14, 2015

M. Smith
L. Kreeger
Cisco Systems, Inc.
November 10, 2014

VXLAN Group Policy Option
draft-smith-vxlan-group-policy-00

Abstract

This document defines a backward compatible extension to Virtual eXtensible Local Area Network (VXLAN) that allows a Tenant System Interface (TSI) Group Identifier to be carried for the purposes of policy enforcement.

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 May 14, 2015.

Copyright Notice

Copyright (c) 2014 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
1.1.	Requirements Language	2
1.2.	Definition of Terms	3
2.	Approach	3
2.1.	VXLAN Group Based Policy Extension	3
3.	Backward Compatibility	4
4.	IANA Considerations	4
5.	Security Considerations	4
6.	Acknowledgements	5
7.	References	5
7.1.	Normative References	5
7.2.	Informative References	5
	Authors' Addresses	5

[1.](#) Introduction

The Group Based Policy [[GROUPPOLICY](#)][GROUPBASEDPOLICY] model defines an application-centric policy model where the application connectivity requirements are specified in a manner that is independent of the underlying network topology. In this model, Tenant System Interfaces (TSIs) are assigned to Tenant System Interface (TSI) Groups. Each TSI Group consists of TSIs that share the same network policies and requirements. Network policies are defined between the TSI Group of the traffic source and the TSI Group of the traffic destination. These policies are deployed when the TSI attaches to the network.

In many situations, the TSI to TSI Group mapping is known only at the Network Virtualization Edge (NVE) that the TSI is attached. This implies that the TSI Group of a packet destination may not be known until the packet reaches the egress NVE where the packet destination is attached. In such situations, it is critical to retain the source TSI Group membership with the packet so that policy can be applied at the egress NVE.

This document defines a backward compatible extension to VXLAN [[RFC7348](#)] that allows the source TSI Group identifier to be carried so that policy can be applied when the destination TSI Group is determined at the egress NVE.

[1.1.](#) 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 [RFC 2119](#) [[RFC2119](#)].

1.2. Definition of Terms

This document uses the same terminology as [\[RFC7365\]](#) and [\[RFC7348\]](#). In addition, the following terms are used:

Tenant System Interface (TSI) Group: A TSI Group is a collection of TSIs that share the same network policies and requirements.

2. Approach

2.1. VXLAN Group Based Policy Extension

The VXLAN Group Based Policy Extension (VXLAN-GBP) header is defined as:




Figure 1: VXLAN-GBP Extension

The following bits are defined in addition to the existing VXLAN fields:

G Bit: Bit 0 of the initial word is defined as the G (Group Based Policy Extension) bit.

G = 1 indicates that the source TSI Group membership is being carried within the Group Policy ID field as defined in this document.

G = 0 indicates that the Group Policy ID is not being carried, and the G Bit MUST be set to 0 as specified in [\[RFC7348\]](#).

D bit: Bit 9 of the initial word is defined as the Don't Learn bit. When set, this bit indicates that the egress VTEP MUST NOT learn the source address of the encapsulated frame.

A Bit: Bit 12 of the initial word is defined as the A (Policy Applied) bit. This bit is only defined as the A bit when the G bit is set to 1.

A = 1 indicates that the group policy has already been applied to this packet. Policies MUST NOT be applied by devices when the A bit is set.

A = 0 indicates that the group policy has not been applied to this packet. Group policies MUST be applied by devices when the A bit is set to 0 and the destination Group has been determined. Devices that apply the Group policy MUST set the A bit to 1 after the policy has been applied.

Group Policy ID: 16 bit identifier that indicates the source TSI Group membership being encapsulated by VXLAN. The allocation of Group Policy ID values is outside the scope of this document.

3. Backward Compatibility

VXLAN [[RFC7348](#)] requires reserved fields to be set to zero on transmit and ignored on receive. This ensures that the G bit will never be set by VXLAN VTEPs and therefore packets received from these VTEPs can be assigned to a default Group Policy ID. It also ensures that VXLAN VTEPs receiving packets with the G bit set will ignore the Group Policy ID. Due to this defined behavior by VXLAN VTEPs, it allows the extensions described in this document to operate on the IANA assigned VXLAN UDP port (port 4789).

In some environments, there may be a mix of devices supporting the VXLAN Group Based Policy Extension and devices that do not. Devices supporting the VXLAN Group Based Policy Extension SHOULD assign traffic arriving without the G bit set to a default Group Policy ID for the purposes of policy enforcement.

4. IANA Considerations

This memo includes no request to IANA.

5. Security Considerations

This document describes an extension to VXLAN to carry the Group Policy Identifier of the source endpoint. These identifiers must be distributed to participating VTEPs that are encapsulating traffic from the endpoints sourcing traffic. While the control plane protocols for distributing these identifiers is outside the scope of this document, any control plane protocol should ensure that these identifiers are securely distributed to the network elements participating in the policy enforcement domain.

Additionally, the Group Policy Identifier field being carried in the packet directly impacts the network policy applied to the traffic.

There is a risk that these identifiers may be spoofed and proper integrity protection should be put in place to ensure that these fields can only be populated by trusted entities. Due to the importance of these fields, confidentiality may also be required to ensure that traffic cannot be targeted for attack based on the policy identifiers. In some environments, these attacks are mitigated through physical security. In other environments, traditional security mechanisms like IPsec that authenticate and optionally encrypt VXLAN traffic including the bits and fields described in this document.

6. Acknowledgements

Many thanks to Tom Edsall and Thomas Graf for their comments and review of this document.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC7348] Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and C. Wright, "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", [RFC 7348](#), August 2014.
- [RFC7365] Lasserre, M., Balus, F., Morin, T., Bitar, N., and Y. Rekhter, "Framework for Data Center (DC) Network Virtualization", [RFC 7365](#), October 2014.

7.2. Informative References

- [GROUPBASEDPOLICY] OpenStack, "Group Based Policy", 2014, <<https://wiki.openstack.org/wiki/GroupBasedPolicy>>.
- [GROUPPOLICY] OpenDaylight, "Group Policy", 2014, <https://wiki.opendaylight.org/view/Group_Policy:Main>.

Authors' Addresses

Michael Smith
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, California 95134
USA

Email: michsmit@cisco.com

Lawrence Kreeger
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, California 95134
USA

Email: kreeger@cisco.com

