Network Working Group
Internet-Draft
Expires: May 14, 2015

F. Xia
B. Sarikaya
Huawei Technologies Co., Ltd.
November 10, 2014

Quality of Service Marking in Virtual eXtensible Local Area Network
draft-xia-nvo3-vxlan-qosmarking-01.txt

Abstract

The Virtual eXtensible Local Area Network enables multiple tenants to operate in a data center. Each tenant needs to be assigned a priority group to prioritize their traffic. Cloud carriers wish to use quality of service to differentiate different applications. For these purposes, six bits are assigned in the eXtensible Local Area Network header. How these bits are assigned and are processed in the network are explained in detail.

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.

Table of Contents

1. Introduction

Data center networks are being increasingly used by telecom operators as well as by enterprises. An important requirement in data center networks is multitenancy, i.e. multiple tenants each with their own isolated network domain. Virtual eXtensible Local Area Network (VXLAN) is a solution that is gaining polularity in industry [RFC7348]. VXLAN overlays a Layer 2 network over a Layer 3 network. Each overlay is identified by the VXLAN Network Identifier (VNI). VXLAN tunnel end point (VTEP) can be hosted at the the hypervisor on the server or higher above in the network. VXLAN encapsulation with a UDP header is only known to the VTEP, the Virtual Machines (VM) never sees it.

It should be noted that in this document, VTEP plays the role of the Network Virtualization Edge (NVE) according to NVO3 architecture for overlay networks like VXLAN or NVGRE defined in [I-D.ietf-nvo3-arch]. NVE interfaces the tenant system underneath with the L3 network called the Virtual Network (VN).

Since VXLAN allows multiple tenants to operate data center operators are facing the problem of treating their traffic. There is interest to provide different quality of service to the tenants based on their service level agreements.

Cloud carriers have interest in different quality of service to different applications such as voice, video, network control applications, etc. In this case, quality of service marking can be

done using deep packet inspection (DPI) in order to detect the type of application in each packet.

In this document, we develop Quality of Service marking solution for VXLAN. The solution is compatible with IP level Differentiated Services model or diffserv described in [RFC2474] and [RFC2475]. Configuration guidelines are described in [RFC4594]. Diffserv interconnection classes and interconnection practice are described in [I-D.geib-tsvwg-diffserv-intercon].

2. Terminology

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]. The terminology in this document is based on the definitions in [RFC7348].

3. Problem Statement

In a VXLAN network multiple tenants are supported. There is interest in assigning different priority to each tenant's traffic based on the premium that tenant paies, etc. In another words, cloud carriers would like to categorize tenants into different traffic classes such as diamond, gold, silver and bronze classes.

Cloud carriers wish to categorize the traffic based on the application such as voice, video, etc. Based on the type of the application different traffic classes may be identified and different priority levels can be assigned to each.

In order to do these, quality of service marking is needed in VXLAN.

The solution proposed in this document is based on using VXLAN header to mark by Network Virtualization Edge (NVE) when the frames are introduced by the virtual machines.

4. QoS Bits in VXLAN Header

Six bits are reserved in VXLAN header flags field shown as QoS-flag in Figure 1.

6 bits called QoS-flag are reserved to indicate the quality of service class that this packet belongs. These bits will be assigned according to the type of traffic carried in this flow, e.g. video, voice, critical application, etc. These assignments will be made adopting IP level Differentiated Services model, diff serv bits or DS field, see Section 7.

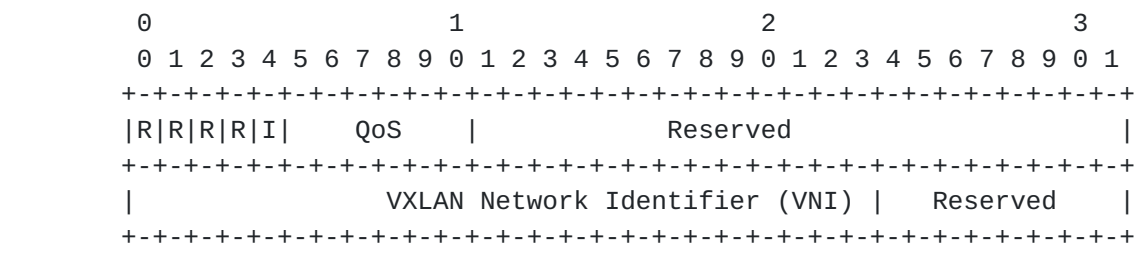

Figure 1: QoS Flag in VXLAN Header

The first three bits (bits 5-7) are precedence bits. They are assigned according to [RFC0791]. Precedence values '110' and '111' are selected for routing traffic.

The last three bits (bits 8-10) are class selector bits. Thet are assigned as follows:

001 - BK or background traffic

000 - BE or best effort traffic

010 - EE or Excellent Effort

011 - CA or Critical Applications

100 - VI or Video

101 - VO or Voice

110 - IC or Internetwork Control

111 - NC or Network Control

'111' has the highest priority while '001' has the lowest, for example, video traffic has higher priority than web surfing which is best effort traffic.

As can be seen the markings are the same as in IEEE 802.1p [IEEE802.1D] which is supported by most switches currently deployed that have the QoS capabilities.

Bits 5,6 and 7 previously reserved are used to assign precedence. Bits 8, 9 and 10 previously reserved are used to assign the class selector bits. The sender SHOULD assign bits 8-10 with bits assigned values as above if the quality of service treatment is needed on this packet. The sender SHOULD assign the same bit pattern to all the

packets of the same flow. The sender MUST assign all other reserved bits to zero.

In real deployment, there are two different mappings to make use of the class selector QoS field.

The first one is based on application priorities. NVE uses some mechanism such as Deep Packet Inspection (DPI) to identify application types, and fills in the QoS field of VXLAN encapsulation based on the identified application types. The below is a possible mapping.

001 - Reserved

000 - ftp/email

010 - web surfing

011 - instant Message

100 - video

101 - voice

110 - High Performance computation

111 - Reserved

The second one is based on tenancy priorities. A cloud carrier could exploit the QoS bits in another different way. The cloud carrier categorizes its tenants into different groups such as diamond, gold, silver, bronze, standard and so on. All traffic for a diamond tenant has a high priority to be forwarded regardless of application types. The below is a possible mapping option.

001 - Reserved

000 - Standard

010 - Bronze

011 - Silver

100 - Gold

101 - Diamond

110 - Emergency

111 - Reserved

5. Quality of Service Operation at VXLAN Decapsulation Point

There are two types of VXLAN packets receivers, that is, a VXLAN enabled NVE or a VXLAN gateway [I-D.sarikaya-nvo3-proxy-vxlan].

When the VXLAN enabled NVE receives the packet, it decapsulates the packet and delivers it downstream to a corresponding VM. If there are multiple packets to be processed, packets with high priority (that is higher QoS value) should be processed first.

The QoS operation is different for the VXLAN gateway processing. The gateway which provides VXLAN tunnel termination functions could be ToR/access switches or switches higher up in the data center network topology. For incoming frames on the VXLAN connected interface, the gateway strips out the VXLAN header and forwards to a physical port based on the destination MAC address of the inner Ethernet frame. If inner VLAN is included in the VXLAN frame or a VLAN is supposed to be added based on configuration, the VXLAN gateway decapsulates the VXLAN packet and remarks the QoS field of the outgoing Ethernet frame based on VXLAN class selector QoS bits. The switch SHOULD copy the class selector Q-Flags of VXLAN encapsulation into IEEE 802.1p Priory Code Point (PCP) field in VLAN tag.

6. Quality of Service Operation at VXLAN encapsulation point

There are two types of VXLAN packet senders, that is, a VXLAN enabled NVE or a VXLAN gateway.

For a VXLAN enabled NVE, the upstream procedure is:

Reception of Frames

The VXLAN enabled NVE receives an Ethernet packet from a hosting VM.

Lookup

Making use of the destination of the Ethernet packet, the VXLAN enabled NVE looks up MAC-NVE mapping table, and retrieves IP address of destination NVE.

Acquisition of QoS parameters

There are two different ways to acquire QoS parameters for VXLAN encapsulation. The first is a dynamic one which requires a VXLAN enabled NVE has Deep Packet Inspection (DPI) capability and can

identify different application types. The second is a static one which requires a VM manager to assign QoS parameters to different VNIs based on premium that different tenancies pay.

Encapsulation of frames

The NVE then encapsulates the packet using VXLAN format with acquired QoS parameters and VNI. The specific format is given in Section 4. After the frame is encapsulated it is sent out upstream to the network.

For a VXLAN gateway, packets are encapsulated using VXLAN format with QoS field in a similar way. Once the VXLAN gateway receives a packet from a non-VXLAN domain, it encapsulates the packet with QoS parameters which are acquired through DPI or priorities of tenancies.

7. QoS processing for VXLAN outer IP header

QoS is user experience of end-to-end network operation. A packet from VM A to VM B normally traverses such network entities sequentially as virtual switch A which is co-located with VM A, TOR switch A, aggregation switch A, a core switch, aggregation switch B, TOR switch B, virtual switch B. VXLAN processing only takes place in virtual switches, and all other network entities only execute IP forwarding. VXLAN QoS mapping to outer IP header at virtual switch A is needed to achieve end-to-end QoS.

Six bits of the Differentiated Services Field (DS field) are used as a codepoint (DSCP) to select the per hop behaviour (PHB) a packet experiences at each node in a Differentiated Services Domain [RFC2474]. DS field is 8 bits long, 6 bits of it are used as DSCP and two bits are unused. DS field is carried in both IPv4 and IPv6 packet headers. The first three bits of DS field are used for IP precedence and the last three are used as diff serv bits. VXLAN outer IP header's DSCP field SHOULD be copied from VXLAN header QoS bits.

Similarly, when a packet forwarded from non-VXLAN domain to VXLAN domain through a VXLAN gateway, DSCP field of outer IP header should be marked based on VXLAN QoS.

8. Security Considerations

Special security considerations in [RFC7348] are applicable.

9. IANA considerations

IANA is requested to assign the Q-Flags bits in VXLAN reserved bits in the header.

10. Acknowledgements

The authors are grateful to David Black and Brian Carpenter for their constructive comments on our work.

11. References

11.1. Normative References

[RFC0826] Plummer, D., "Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware", STD 37, RFC 826, November 1982.

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.

[RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, December 1998.

[RFC2597] Heinanen, J., Baker, F., Weiss, W., and J. Wroclawski, "Assured Forwarding PHB Group", RFC 2597, June 1999.

[RFC4594] Babiarz, J., Chan, K., and F. Baker, "Configuration Guidelines for DiffServ Service Classes", RFC 4594, August 2006.

[I-D.ietf-nvo3-arch] Black, D., Hudson, J., Kreeger, L., Lasserre, M., and T. Narten, "An Architecture for Overlay Networks (NVO3)", draft-ietf-nvo3-arch-02 (work in progress), October 2014.

[IEEE802.1D]
IEEE, "Virtual Bridged Local Area Networks", IEEE Std 802.1D-2005, May 2006.

11.2. Informative References

[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.

[I-D.geib-tsvwg-diffserv-intercon]
Geib, R. and D. Black, "DiffServ interconnection classes and practice", draft-geib-tsvwg-diffserv-intercon-07 (work in progress), October 2014.

[I-D.sarikaya-nvo3-proxy-vxlan]
Sarikaya, B. and F. Xia, "Virtual eXtensible Local Area Network over IEEE 802.1Qbg", draft-sarikaya-nvo3-proxy-vxlan-00 (work in progress), October 2014.

Authors' Addresses

Frank Xia
Huawei Technologies Co., Ltd.
101 Software Avenue, Yuhua District
Nanjing, Jiangsu 210012, China

Phone: ++86-25-56625443
Email: xiayangsong@huawei.com

Behcet Sarikaya
Huawei Technologies Co., Ltd.
5340 Legacy Dr. Building 3
Plano, TX 75024

Phone: +1 972-509-5599
Email: sarikaya@ieee.org