

NVO3 WG
Internet-Draft
Intended status: Standards Track
Expires: April 14, 2016

Fangwei. Hu
Ran. Chen
ZTE Corporation
Mallik. Mahalingam
Springpath
Qiang. Zu
Ericsson
October 12, 2015

YANG Data Model for VxLAN Protocol
draft-chen-nvo3-vxlan-yang-01.txt

Abstract

This document defines a YANG data model for VxLAN protocol.

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 April 14, 2016.

Copyright Notice

Copyright (c) 2015 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. Terminology	2
3. YANG Data Model for VxLAN Configuration	2
3.1. VxLAN Multicast IP Address	2
3.2. VxLAN Mapping Type	2
3.3. VTEP Mode	3
4. Design Tree of Configuration Data Model	3
5. VxLAN YANG Configuration Data Model	4
6. Design Tree of Operational State Model	9
7. VxLAN YANG Operational State Model	10
8. Security Considerations	13
9. Acknowledgements	13
10. IANA Considerations	13
11. Normative References	13
Authors' Addresses	14

1. Introduction

YANG[RFC6020] is a data definition language that was introduced to define the contents of a conceptual data store that allows networked devices to be managed using NETCONF [RFC6241]. This document defines a YANG data model for the configuration of VxLAN protocol [RFC7348].

2. Terminology

3. YANG Data Model for VxLAN Configuration

3.1. VxLAN Multicast IP Address

The `vxlan-multicast-ip` is used to configure the IP multicast group, which the VxLAN VNI of the VTEP is mapping to.

3.2. VxLAN Mapping Type

The `VxLAN-mapping-type` is used to configure the VxLAN mapping type for the interface. The value is as following:

- o `vlan-1:1`: the vxlan mapping type is VLAN, and each VxLAN is only mapping to one VLAN.
- o `vlan- n:1`: the vxlan mapping type is VLAN, and each VxLAN could be mapped to several VLANs.
- o `L3-interface`: the VxLAN mapping type is layer 3 interface.
- o `mac`: the VxLAN mapping type is MAC address.

- o vlan-l2-interface: the VxLAN mapping type is vlan + Layer 2 interface.

3.3. VTEP Mode

The VTEP-mode is used to configure the mode for the VTEP interface. There are two modes: discard-inner-vlan and no-discard-inner-vlan. If the VTEP interface works in the discard-inner-vlan mode, the VxLAN is only mapping to one VLAN. The inner VLAN tag will be stripped when encapsulating the VxLAN frame. On the decapsulation side, if VTEP receives the VxLAN frame with inner VLAN tag, it will discard the frame in this work mode. If the VTEP receives the VxLAN frame without VLAN tag, it will fill in the VLAN tag based on the VxLAN and VLAN mapping.

If the VTEP interface works in the no-discard-inner-vlan mode, the VxLAN could be mapped to several VLANs. The inner VLAN tag will not be stripped when encapsulating the VxLAN frame in the VxLAN encapsulation side. On the decapsulation side, if VTEP receives the VxLAN frame, it will strip the VxLAN header, and keep the VLAN frame.

4. Design Tree of Configuration Data Model

```

module: ietf-vxlan
  +--rw vxlan
    +--rw vxlan-instance* [vxlan-id]
      +--rw vxlan-id          uint32
      +--rw multicast-ip      inet:ipv4-address
      +--rw (vxlan-map-types)?
        +--:(vxlan-map-vlan)
          +--rw map-type-vlan? enumeration
          +--rw map-vlan* [vlan-id]
            +--rw vlan-id    uint16
        +--:(vxlan-map-mac)
          +--rw map-type-mac? empty
          +--rw mac          yang:mac-address
        +--:(vxlan-map-l2interface)
          +--rw map-type-l2interface? empty
          +--rw vlan-id          uint16
          +--rw interface-name    if:interface-ref
        +--:(vxlan-map-l3interface)
          +--rw map-type-l3interface? empty
          +--rw map-l3interface* [interface-name]
            +--rw interface-name    if:interface-ref
    +--rw interfaces
      +--rw interface* [name]
        +--rw name          if:interface-ref
        +--rw vtep-instances* [vtep-id]
          +--rw vtep-id      uint32
          +--rw vtep-name?   string
          +--rw source-interface? if:interface-ref
          +--rw vtep-mode?   enumeration
          +--rw bind-vxlan-id* [vxlan-id]
            +--rw vxlan-id    uint32

```

5. VxLAN YANG Configuration Data Model

```

module ietf-vxlan {
  namespace "urn:ietf:params:xml:ns:yang:ietf-vxlan";
  prefix "vxlan";

  import ietf-interfaces {
    prefix "if";
  }

  import ietf-inet-types {
    prefix "inet";
  }
}

```

```
import ietf-yang-types {
  prefix yang;
}

organization
  "IETF NVO3(Network Virtualization Overlays) Working Group";

contact
  "
    WG List: <mailto:nvo3@ietf.org>

    WG Chair: Matthew Bocci
              <mailto:matthew.bocci@alcatel-lucent.com>

    WG Chair: Benson Schliesser
              <mailto:bensons@queuefull.net>

    Editor:   Ran Chen
              <mailto:chen.ran@zte.com.cn>

    Editor:   Fangwei Hu
              <mailto:hu.fangwei@zte.com.cn>

    Editor:   Mallik Mahalingam
              <mailto:mallik_mahalingam@yahoo.com>

    Editor:   Zu Qiang
              <mailto:Zu.Qiang@Ericsson.com>
  ";

description "The YANG module defines a generic configuration
  model for VxLAN protocol";

revision 2015-10-12 {
  description
    "01 revision.";
  reference
    "draft-chen-nvo3-vxlan-yang-01";
}

revision 2015-05-05 {
  description "Initial revision";
  reference
    "draft-chen-nvo3-vxlan-yang-00";
}

/* Configuration Data */
container vxlan{
```

```
list vxlan-instance {
  key vxlan-id ;
  leaf vxlan-id {
    type uint32 {
      range 1..16777215 {
        description 'The value of VXLAN,it must
          between 1 to 16777215.' ;
      }
    }
    description 'VXLAN ID.' ;
  }

  leaf multicast-ip {
    type inet:ipv4-address ;
    mandatory true ;
    description 'VXLAN multicast IP address.' ;
  }

  choice vxlan-map-types {
    case vxlan-map-vlan {
      leaf map-type-vlan {
        type enumeration {
          enum map-type-vlanl1t1 {
            value 1 ;
            description 'Map type VLAN1:1.' ;
          }
          enum map-type-vlanl1n {
            value 2 ;
            description 'Map type VLAN1:n.' ;
          }
        }
        default map-type-vlanl1t1 ;
        description 'VXLAN map VLAN.' ;
      }

      list map-vlan {
        key vlan-id ;
        leaf vlan-id {
          type uint16 {
            range 1..4094 {
              description 'The value of VLAN,it must
                between 1 to 4094.' ;
            }
          }
        }
      }
    }
  }
}
```

```
        description 'VLAN ID.' ;
    }
    description 'VXLAN map VLAN.' ;
}
description 'VXLAN map VLAN.' ;
}

case vxlan-map-mac {
    leaf map-type-mac {
        type empty ;
        description 'VXLAN map MAC.' ;
    }

    leaf mac {
        type yang:mac-address ;
        mandatory true ;
        description 'MAC address.' ;
    }
    description 'VXLAN map MAC.' ;
}

case vxlan-map-l2interface {
    leaf map-type-l2interface {
        type empty ;
        description 'VXLAN map L2 interface.' ;
    }
    leaf vlan-id {
        type uint16 {
            range 1..4094 {
                description 'The value of VLAN,it must
                between 1 to 4094.' ;
            }
        }
        mandatory true ;
        description 'VLAN ID.' ;
    }

    leaf interface-name {
        type if:interface-ref;
        mandatory true ;
        description 'L2 interface name.' ;
    }
    description 'VXLAN map L2 interface.' ;
}

case vxlan-map-l3interface {
    leaf map-type-l3interface {
        type empty ;
```

```
        description 'VXLAN map L3 interface.' ;
    }

    list map-l3interface {
        key interface-name ;
        leaf interface-name {
            type if:interface-ref;
            description 'L3 interface name.' ;
        }
        description 'VXLAN map L3 interface.' ;
    }
    description 'VXLAN map L3 interface.' ;
}
description 'VXLAN map type choice.' ;
}
description 'VXLAN instance.' ;
}

container interfaces {
    list interface{
        key "name";
        leaf name {
            type if:interface-ref;
            description 'vtep interface name';
        }
    }
    list vtep-instances {
        key vtep-id ;
        leaf vtep-id {
            type uint32;
            description 'VTEP id.' ;
        }
    }

    leaf vtep-name{
        type string;
        description 'VTEP instance name.' ;
    }

    leaf source-interface {
        type if:interface-ref;
        description 'Source interface name.' ;
    }

    leaf vtep-mode {
        type enumeration {
            enum discard-inner-vlan {
                value 1 ;
                description 'Discard inner-VLAN.' ;
            }
        }
    }
}
```

```
        enum no-discard-inner-vlan {
            value 2 ;
            description 'No discard inner-VLAN.' ;
        }
    }
    default discard-inner-vlan ;
    description 'VTEP mode.' ;
}

list bind-vxlan-id {
    key vxlan-id;
    leaf vxlan-id {
        type uint32;
        description 'VXLAN ID.' ;
    }
    description 'the vxlan id list for the vtep.';
}
description 'the vtep instance.';
}
description 'vtep interface.';
}
description 'the vtep interfaces configuration.';
}
description 'the vxlan configure model.';
}
```

6. Design Tree of Operational State Model

```

+--ro vxlan-state
  +--ro vxlan-instance
    |   +--ro vxlan-id?                uint32
    |   +--ro multicast-ip?           inet:ipv4-address
    |   +--ro (vxlan-map-types)?
    |     +---:(vxlan-map-vlan)
    |       |   +--ro map-type-vlan?   enumeration
    |       |   +--ro map-vlan* [vlan-id]
    |       |     +--ro vlan-id      uint16
    |       +---:(vxlan-map-mac)
    |         |   +--ro map-type-mac?   empty
    |         |   +--ro mac?          yang:mac-address
    |       +---:(vxlan-map-l2interface)
    |         |   +--ro map-type-l2interface? empty
    |         |   +--ro vlan-id?      uint16
    |         |   +--ro interface-name? if:interface-ref
    |       +---:(vxlan-map-l3interface)
    |         |   +--ro map-type-l3interface? empty
    |         |   +--ro map-l3interface* [interface-name]
    |         |     +--ro interface-name if:interface-ref
    +--ro vtep-instance
      +--ro vtep-id?                uint32
      +--ro vtep-name?              string
      +--ro source-interface?       if:interface-ref
      +--ro vtep-mode?              enumeration
      +--ro bind-vxlan-id* [vxlan-id]
      +--ro vxlan-id                uint32

```

7. VxLAN YANG Operational State Model

```

/* Operational data */
container vxlan-state{
  container vxlan-instance {
    leaf vxlan-id {
      type uint32;
      description "show the vxlan id.";
    }
    leaf multicast-ip {
      type inet:ipv4-address ;
      description 'show VxLAN multicast IP address.' ;
    }
  }
  choice vxlan-map-types {
    case vxlan-map-vlan {
      leaf map-type-vlan {

```

```
    type enumeration {
      enum map-type-vlanlto1 {
        value 1 ;
        description 'Map type VLAN1:1.' ;
      }
      enum map-type-vlanlton {
        value 2 ;
        description 'Map type VLAN1:n.' ;
      }
    }
    default map-type-vlanlto1 ;
    description 'VXLAN map VLAN.' ;
  }

  list map-vlan {
    key vlan-id ;
    leaf vlan-id {
      type uint16;
      description 'VLAN ID.' ;
    }
    description 'VXLAN map VLAN.' ;
  }
  description 'VXLAN map VLAN.' ;
}

case vxlan-map-mac {
  leaf map-type-mac {
    type empty ;
    description 'VXLAN map MAC.' ;
  }

  leaf mac {
    type yang:mac-address ;
    description 'MAC address.' ;
  }
  description 'VXLAN map MAC.' ;
}

case vxlan-map-l2interface {
  leaf map-type-l2interface {
    type empty ;
    description 'VXLAN map L2 interface.' ;
  }
  leaf vlan-id {
    type uint16;
    description 'VLAN ID.' ;
  }
}
```

```
    leaf interface-name {
      type if:interface-ref;
      description 'L2 interface name.' ;
    }
    description 'VXLAN map L2 interface.' ;
  }

  case vxlan-map-l3interface {
    leaf map-type-l3interface {
      type empty ;
      description 'VXLAN map L3 interface.' ;
    }

    list map-l3interface {
      key interface-name ;
      leaf interface-name {
        type if:interface-ref;
        description 'L3 interface name.' ;
      }
      description 'VXLAN map L3 interface.' ;
    }
    description 'VXLAN map L3 interface.' ;
  }
  description 'VXLAN map type choice.' ;
}
description 'show the vxlan instance information.' ;
}

container vtep-instance {
  leaf vtep-id {
    type uint32;
    description "show the vtep id.";
  }
  leaf vtep-name {
    type string;
    description 'show the vtep name.' ;
  }
  leaf source-interface {
    type if:interface-ref;
    description 'show the source interface.' ;
  }

  leaf vtep-mode {
    type enumeration {
      enum discard-inner-vlan {
        value 1 ;
        description 'Discard inner-VLAN.' ;
      }
    }
  }
}
```

```
        enum no-discard-inner-vlan {
            value 2 ;
            description 'No discard inner-VLAN.' ;
        }
    }
    default discard-inner-vlan ;
    description 'show VTEP mode.' ;
}

list bind-vxlan-id {
    key vxlan-id;
    leaf vxlan-id {
        type uint32;
        description 'show the VXLAN ID.' ;
    }
    description 'show the vxlan id list for the vtep.';
}
description 'show the vtep information';
}
description 'show the vxlan operational state information';
}
```

8. Security Considerations

9. Acknowledgements

10. IANA Considerations

This document requires no IANA Actions. Please remove this section before RFC publication.

11. Normative References

- [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, DOI 10.17487/RFC6020, October 2010, <<http://www.rfc-editor.org/info/rfc6020>>.
- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<http://www.rfc-editor.org/info/rfc6241>>.
- [RFC6991] Schoenwaelder, J., Ed., "Common YANG Data Types", RFC 6991, DOI 10.17487/RFC6991, July 2013, <<http://www.rfc-editor.org/info/rfc6991>>.

- [RFC7223] Bjorklund, M., "A YANG Data Model for Interface Management", RFC 7223, DOI 10.17487/RFC7223, May 2014, <<http://www.rfc-editor.org/info/rfc7223>>.
- [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, DOI 10.17487/RFC7348, August 2014, <<http://www.rfc-editor.org/info/rfc7348>>.

Authors' Addresses

Fangwei Hu
ZTE Corporation
No.889 Bibo Rd
Shanghai 201203
China

Phone: +86 21 68896273
Email: hu.fangwei@zte.com.cn

Ran Chen
ZTE Corporation
No.50 Software Avenue, Yuhuatai District
Nanjing, Jiangsu Province 210012
China

Phone: +86 025 88014636
Email: chen.ran@zte.com.cn

Mallik Mahalingam
Springpath
640 W. California Ave, Suite #110
Sunnyvale, CA 94086
USA

Email: mallik_mahalingam@yahoo.com

Zu Qiang
Ericsson
8400, boul. Decarie
Ville Mont-Royal, QC
Canada

Email: Zu.Qiang@Ericsson.com