

INTERNET-DRAFT
Intended Status: Standards Track

Mingui Zhang
Lianshu Zheng
Feng Zheng
Huawei
Fu Qiao
China Mobile
October 21, 2014

Expires: April 24, 2015

YANG Data Model for NV03 Protocols
draft-zhang-nvo3-yang-cfg-00.txt

Abstract

This document describes the base YANG data model that can be used by operators to configure and manage NV03 protocols.

Status of this Memo

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

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

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

Copyright and License 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

INTERNET-DRAFT YANG Data Model for NV03 Protocols October 21, 2014

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	3
2.	Acronyms and Terminology	3
2.1.	Acronyms	3
2.2.	Terminology	3
3.	The YANG Data Model for NV03	3
3.1.	The Configuration Parameters	3
3.1.1.	NVE ID	4
3.1.2.	Encapsulation Type	4
3.1.3.	Source and Destination UDP Port	4
3.1.4.	Virtual Network Instance	4
3.1.5.	Flags in the Header	4
3.1.6.	BUM Mode	5
3.2.	Statistics	5
3.3.	Model Structure	5
3.4.	YANG Module	6
4.	Security Considerations	13
5.	IANA Considerations	13
	Acknowledgements	13
6.	References	13
6.1.	Normative References	13
6.2.	Informative References	14
	Author's Addresses	15

INTERNET-DRAFT YANG Data Model for NV03 Protocols October 21, 2014

[1.](#) Introduction

Network Virtualization Overlays (NV03) enable network virtualization for data center networks environment that assumes an IP-based underlay.

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 specifies a YANG data model that can be used to configure and manage NV03 protocols. The model covers the configuration of preliminary NV03 instances as well as their operation states. We call it the base model for NV03 in this document.

Two default data plane encapsulation techniques for NV03 are incorporated in the base module. It allows other possible encapsulations developed in the industry later on to be included by introducing a new encapsulation as a new 'choice'. As new solutions for NV03 are developed, corresponding YANG data modules may be specified. The base module can be augmented to accommodate these new solutions.

[2.](#) Acronyms and Terminology

[2.1.](#) Acronyms

NV03: Network Virtualization Overlays

VNI: Virtual Network Instance

BUM: Broadcast, Unknown Unicast, Multicast traffic

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

Familiarity with [[RFC7364](#)], [[RFC7365](#)], [I-D.ietf-nvo3-dataplane-requirements] and [[RFC7348](#)] is assumed in this document.

[3.](#) The YANG Data Model for NV03

The NV03 base YANG module is divided in three containers. The first container contains writable parameters. The second container contains the writable enablers per VNI for the statistical operational states as well as the status of these enablers. The third container contains the statistical operational states.

[3.1.](#) The Configuration Parameters

[3.1.1.](#) NVE ID

The 'nvo3Nves' list contains NVEs under the management. The NVE is identified using the 'srcAddr', which is the underlay IP address of the NVE.

[3.1.2.](#) Encapsulation Type

The encapsulation type of NV03 is to be determined by the operators according to their network. NV03 encapsulation arising from the industry is not unique. The value of 'encapType' determines the encapsulation. Besides the default VxLAN encapsulation [[RFC7348](#)], the module also allows other possible encapsulation, for example the NVGRE encapsulation [[I-D.sridharan-virtualization-nvgre](#)]. It can be extended to include more encapsulation choices.

[3.1.3.](#) Source and Destination UDP Port

If the encapsulation type is VxLAN, it's recommended the source UDP port is calculated using a hash of fields (e.g., source IP address, destination IP address, protocol type, source MAC address and destination MAC address) from the inner packet. The bit value of 'srcUdpPortGenRuleFlags' indicates whether each field is used in the hashing function with the high-order bit indicating the first hashing field.

The value of destination UDP port is set to be 'destUdpPort'. For VxLAN, IANA has assigned the value 4789 for the UDP destination port, and this value SHOULD be used by default as the destination UDP port

[[RFC7348](#)].

[3.1.4.](#) Virtual Network Instance

A Virtual Network Instance ('VNI') is a specific VN instance on an NVE [[RFC7365](#)]. At each NVE, a Tenant System is connect to VNIs through Virtual Access Points (VAP). VAPs can be physical ports or virtual ports identified by the bridge domain Identifier ('bdId'). The mapping between VNI and bdId is managed by the operator.

[3.1.5.](#) Flags in the Header

Flags in the NV03 header are to be configured. For VxLAN, the I bit of 'flag' MUST be set to 1 while other 7 bits are reserved and MUST be set to zero on transmission and ignored on receipt. For NVGRE, bits in position 0 and 3 MUST be set to zero. The bit in position 2 MUST be set to 1. Bits from position 4 through 12 are reserved and MUST be set to zero on transmission and ignored on receipt.

[3.1.6.](#) BUM Mode

An NVE SHOULD support either ingress replication or point to multipoint tunnels [[I-D.ietf-nvo3-dataplane-requirements](#)] on a per-VNI basis. It is possible that both modes be used simultaneously in one NV03 network by different NVEs.

If ingress replication is used, the receiver addresses are listed in 'peerAddr'. If the choice is point to multipoint tunnels, the multicast address is given as 'multiAddr'.

[3.2.](#) Statistics

Operators can determine whether a NVE should gather statistic values on a per-VNI base. The enablers are contained in the 'nvo3Info' list as 'statisticsEnable' leaf.

If the gathering for a VNI is enabled, the statistical information about the local NVEs, the remote NVEs, the flows and the MAC addresses will be collected by the NVEs in this VNI.

[3.3.](#) Model Structure

```

module: nvo3-base
  +--rw nvo3Nves
  |   +--rw nvo3Nve* [ifName]
  |       +--rw ifName                string
  |       +--rw srcAddr?              inet:ip-address
  |       +--rw (encapType)?
  |       |   +--:(vxlan)
  |       |       +--rw srcUdpPortGenRule?  uint8
  |       |       +--rw destUdpPort?      uint16
  |       +--rw members
  |       |   +--rw member* [VNI]
  |       |       +--rw VNI                uint32
  |       |       +--rw bdId                uint32
  |       |       +--rw (encapType)?
  |       |       |   +--:(vxlan)
  |       |       |   |   +--rw vxlanFlag?  flags-vxlan
  |       |       |   |   +--:(nvgre)
  |       |       |   |       +--rw nvgreFlag?  flags-nvgre
  |       |       |   |       +--rw flowId?    uint8
  |       |       +--rw (bumMode)?
  |       |       |   +--:(headEnd)
  |       |       |   |   +--rw peerAddr*    inet:ip-address
  |       |       |   +--:(multiGroup)
  |       |       |       +--rw multiAddr?   inet:ip-address
  |   +--rw nvo3Infos

```

```

  |   +--rw nvo3Info* [VNI]
  |       +--rw VNI                uint32
  |       +--rw statisticsEnable?  enumeration
  |       +--ro status?            enumeration
  +--rw nvo3Statistics
  |   +--ro localNVE*
  |       |   +--ro count?    uint64
  |       |   +--ro ipAddr*   inet:ip-address
  |   +--ro remoteNVE*
  |       |   +--ro count?    uint64
  |       |   +--ro ipAddr*   inet:ip-address
  +--rw flowStatistics
  |   +--ro instanceCount?  uint64
  |   +--ro flowStatistic*
  |       +--ro VNI?                uint32

```

```

|      +---ro inPacketsCount?          uint64
|      +---ro inBytesCount?            uint64
|      +---ro inUnicasts?              uint64
|      +---ro inMulticasts?            uint64
|      +---ro inBroadcasts?            uint64
|      +---ro inUnknownUnicastDrops?   uint64
|      +---ro inUnknownMulticastDrops? uint64
|      +---ro inBroadcastsDrops?       uint64
|      +---ro outPacketsCount?          uint64
|      +---ro outBytesCount?            uint64
|      +---ro outUnicasts?              uint64
|      +---ro outMulticasts?            uint64
|      +---ro outBroadcasts?            uint64
+---rw MacStatistics
    +---ro MacStatistic*
        +---ro VNI?                    uint32
        +---ro vmMacCount?              uint64

```

Figure 3.1. The tree structure of YANG module for NV03 configuration

3.4. YANG Module

```

<CODE BEGIN>
module nvo3-base {
  namespace "http://www.huawei.com/netconf/nvo3";
  //namespace need to be assigned by IANA
  prefix "nvo3";
  import ietf-inet-types {
    prefix "inet";
  }
  organization "IETF NV03 Working Group";
  contact "vero.zheng@huawei.com
    zhangmingui@huawei.com

```

```

    habby.zheng@huawei.com";
  description "nvo3 yang module";
  revision "2014-10-21";

  typedef flags-vxlan {
    type bits {
      bit validVni {
        position 4;

```

```

    }
  }
}

typedef flags-nvgre {
  type bits {
    bit zero{
      position 0;
    }
    bit KeyPresent{
      position 2;
    }
    bit three{
      position 3;
    }
  }
}

container nvo3Nves {
  list nvo3Nve {
    key "ifName";
    leaf ifName {
      description
        "Interface name.";
      type string;
      mandatory "true";
    }
    leaf srcAddr {
      description
        "Local NVE address.";
      type inet:ip-address;
    }
    choice encapType{
      case vxlan{
        leaf srcUdpPortGenRule {
          description
            "The rule of generating source udp port.";
          type uint8;
        }
        leaf destUdpPort {

```



```

        "Destination udp port.";
        default "4789";
        type uint16{
            range "0 .. 65535";
        }
    }
}
}
container members {
    list member {
        key "VNI";
        leaf VNI {
            description
                "Virutal Network Instance (VNI). For VxLAN, it is
                VXLAN Network Identifier; For NVGRE, it is Virtual
                Subnet Identifier.";
            mandatory "true";
            type uint32 {
                range "1 .. 16777215";
            }
        }
        leaf bdId {
            description
                "Bridge Domain ID";
            mandatory "true";
            type uint32 {
                range "1 .. 32768";
            }
        }
    }
    choice encapType{
        case vxlan{
            leaf vxlanFlag {
                description
                    "The 8 bits flags of VXLAN. The I flag MUST be
                    set to 1 for a valid VNI. Others are reserved.";
                type flags-vxlan;
            }
        }
        case nvgre{
            leaf nvgreFlag {
                description
                    "The 13 bits flags of NVGRE. Bit 0 and 3 are set
                    to 0 while bit 2 is set to 1. Others are reserved";
                type flags-nvgre;
            }
            leaf flowId {
                description

```

```

        "This is an 8-bit value that is used to provide
        per-flow entropy for flows in the same VSID.
        If a FlowID is not generated, it MUST be set to
        all zero.";
    type uint8;
}
}
choice bumMode{
    case headEnd{
        leaf-list peerAddr {
            description
                "Remote NVE address.";
            type inet:ip-address;
        }
    }
    case multiGroup{
        leaf multiAddr {
            description
                "The multicast group address.";
            type inet:ip-address;
        }
    }
}
}
}
} //container members
}
} //container nvo3Nves

container nvo3Infos {
    list nvo3Info {
        key "VNI";
        leaf VNI {
            description
                "Virtual Network Instance (VNI). For VxLAN, it is
                VXLAN Network Identifier; For NVGRE, it is Virtual
                Subnet Identifier.";
            mandatory "true";
            type uint32 {
                range "1 .. 16777215";
            }
        }
    }
    leaf statisticsEnable {
        type enumeration {
            enum enable {

```



```

        config "false";
        type inet:ip-address;
    }
}
list remoteNVE {
    config "false";
    leaf count {
        description
            "The value of count.";
    }
}

```

```

        config "false";
        type uint64;
    }
    leaf-list ipAddr {
        description
            "List of remote NVE address.";
        config "false";
        type inet:ip-address;
    }
}
container flowStatistics {
    grouping inputStatistics {
        leaf inPacketsCount {
            config "false";
            type uint64;
        }
        leaf inBytesCount {
            config "false";
            type uint64;
        }
        leaf inUnicasts {
            config "false";
            type uint64;
        }
        leaf inMulticasts {
            config "false";
            type uint64;
        }
        leaf inBroadcasts {
            config "false";
            type uint64;
        }
    }
}

```

```

    leaf inUnknownUnicastDrops {
        config "false";
        type uint64;
    }
    leaf inUnknownMulticastDrops {
        config "false";
        type uint64;
    }
    leaf inBroadcastsDrops {
        config "false";
        type uint64;
    }
}
grouping outputStatistics {
    leaf outPacketsCount {
        config "false";

```

```

        type uint64;
    }
    leaf outBytesCount {
        config "false";
        type uint64;
    }
    leaf outUnicasts {
        config "false";
        type uint64;
    }
    leaf outMulticasts {
        config "false";
        type uint64;
    }
    leaf outBroadcasts {
        config "false";
        type uint64;
    }
}
leaf instanceCount {
    config "false";
    type uint64;
}
list flowStatistic{
    config "false";

```

```

leaf VNI {
  description
    "Virtual Network Instance (VNI). For VxLAN, it is
    VXLAN Network Identifier; For NVGRE, it is Virtual
    Subnet Identifier.";
  config "false";
  type uint32;
}
uses inputStatistics;
uses outputStatistics;
}
} //container flowStatistics
container MacStatistics {
  list MacStatistic{
    config "false";
    leaf VNI {
      description
        "Virtual Network Instance (VNI). For VxLAN, it is
        VXLAN Network Identifier; For NVGRE, it is Virtual
        Subnet Identifier.";
      config "false";
      type uint32;
    }
  }
}

```

```

leaf vmMacCount {
  description
    "Count of learning Mac addresses for the VNI.";
  config "false";
  type uint64;
}
}
} //container MacStatistics
} //container nvo3Statistics
}
<CODE END>

```

[4. Security Considerations](#)

This document raises no new security issues.

[5. IANA Considerations](#)

The namespace URI defined in [Section 3.3](#) need be registered in the IETF XML registry [[RFC3688](#)].

This document need to register the 'nvo3-base' YANG module in the YANG Module Names registry [[RFC6020](#)].

Acknowledgements

Authors would like to thank the comments and suggestions from Tao Han, Weilian Jiang.

[6](#). References

[6.1](#). Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC7364] T. Narten, E. Gray, et al, "Problem Statement: Overlays for Network Virtualization", [draft-ietf-nvo3-overlay-problem-statement](#), working in progress.
- [RFC7365] Marc Lasserre, Florin Balus, et al, "Framework for DC Network Virtualization", [draft-ietf-nvo3-framework](#), working in progress.
- [I-D.ietf-nvo3-dataplane-requirements] Nabil Bitar, Marc Lasserre, et al, "NV03 Data Plane Requirements", [draft-ietf-nvo3-dataplane-requirements](#), working in progress.

- [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.
- [RFC3688] Mealling, M., "The IETF XML Registry", [BCP 81](#), [RFC 3688](#), January 2004.
- [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", [RFC 6020](#),

October 2010.

[RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", [RFC 6241](#), June 2011.

[6.2](#). Informative References

[I-D.zhang-nvo3-yang-active-active] M. Zhang, J. Xia, "YANG Data Model for NVO3 Active Active Access", [draft-zhang-nvo3-yang-active-active](#), working in progress.

[I-D.sridharan-virtualization-nvgre] M. Sridharan, A. Greenberg, et al, "NVGRE: Network Virtualization using Generic Routing Encapsulation", [draft-sridharan-virtualization-nvgre](#), working in progress.

Huawei Technologies
No. 156 Beiqing Rd. Haidian District,
Beijing 100095
P.R. China

EMail: zhangmingui@huawei.com

Lianshu Zheng
Huawei Technologies
No. 156 Beiqing Rd. Haidian District,
Beijing 100095
P.R. China

EMail: vero.zheng@huawei.com

Feng Zheng
Huawei Technologies
No. 156 Beiqing Rd. Haidian District,
Beijing 100095
P.R. China

EMail: habby.zheng@huawei.com

Fu Qiao
China Mobile

EMail: fuqiao@chinamobile.com