

Internet Engineering Task Force  
Ed.  
Internet-Draft  
Ball  
Intended status: Standards Track  
Singh  
Expires: May 7, 2020  
Systems  
  
Sivaraj  
  
Networks  
  
2019

R. Wilton,  
  
D.  
  
T.  
  
Cisco  
  
S.  
  
Juniper  
  
November 4,

**Sub-interface VLAN YANG Data Models**  
**draft-ietf-netmod-sub-intf-vlan-model-06**

Abstract

This document defines YANG modules to add support for classifying traffic received on interfaces as Ethernet/VLAN framed packets to sub-interfaces based on the fields available in the Ethernet/VLAN frame headers. These modules allow configuration of Layer 3 and Layer 2 sub-interfaces (e.g. attachment circuits) that can interoperate with IETF based forwarding protocols; such as IP and L3VPN services; or L2VPN services like VPWS, VPLS, and EVPN. The sub-interfaces also interoperate with VLAN tagged traffic originating from an IEEE 802.1Q compliant bridge.

The model differs from an IEEE 802.1Q bridge model in that the configuration is interface/sub-interface based as opposed to being based on membership of an 802.1Q VLAN bridge.

The YANG data models in this document conforms to the Network Management Datastore Architecture (NMDA) defined in [RFC 8342](#).

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 <https://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 7, 2020.



## Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](https://trustee.ietf.org/license-info) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://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

<a href="#">1.</a>	Introduction . . . . .	3
<a href="#">1.1.</a>	Terminology . . . . .	3
<a href="#">1.2.</a>	Tree Diagrams . . . . .	4
<a href="#">2.</a>	Objectives . . . . .	4
<a href="#">2.1.</a>	Interoperability with IEEE 802.1Q compliant bridges . . . . .	4
<a href="#">3.</a>	L3 Interface VLAN Model . . . . .	4
<a href="#">4.</a>	Flexible Encapsulation Model . . . . .	5
<a href="#">5.</a>	L3 Interface VLAN YANG Module . . . . .	7
<a href="#">6.</a>	Flexible Encapsulation YANG Module . . . . .	10
<a href="#">7.</a>	Examples . . . . .	19
<a href="#">7.1.</a>	Layer 3 sub-interfaces with IPv6 . . . . .	19
<a href="#">7.2.</a>	Layer 2 sub-interfaces with L2VPN . . . . .	21
<a href="#">8.</a>	Acknowledgements . . . . .	23
<a href="#">9.</a>	ChangeLog . . . . .	23
<a href="#">9.1.</a>	WG version -05 . . . . .	24
<a href="#">9.2.</a>	WG version -04 . . . . .	24
<a href="#">9.3.</a>	WG version -03 . . . . .	24

<a href="#">24</a>	<a href="#">9.4.</a>	WG version -02 . . . . .
<a href="#">24</a>	<a href="#">9.5.</a>	WG version -01 . . . . .
<a href="#">24</a>	<a href="#">9.6.</a>	Version -04 . . . . .
<a href="#">24</a>	<a href="#">9.7.</a>	Version -03 . . . . .
<a href="#">25</a>	<a href="#">10.</a>	IANA Considerations . . . . .
<a href="#">25</a>	<a href="#">11.</a>	Security Considerations . . . . .
<a href="#">25</a>	<a href="#">11.1.</a>	if-l3-vlan.yang . . . . .
<a href="#">26</a>	<a href="#">11.2.</a>	flexible-encapsulation.yang . . . . .
<a href="#">27</a>	<a href="#">12.</a>	References . . . . .
<a href="#">27</a>	<a href="#">12.1.</a>	Normative References . . . . .
<a href="#">28</a>	<a href="#">12.2.</a>	Informative References . . . . .
<a href="#">29</a>	<a href="#">Appendix A.</a>	Comparison with the IEEE 802.1Q Configuration Model
<a href="#">29</a>	<a href="#">A.1.</a>	Sub-interface based configuration model overview . . . . .
<a href="#">30</a>	<a href="#">A.2.</a>	IEEE 802.1Q Bridge Configuration Model Overview . . . . .

[A.3](#). Possible Overlap Between the Two Models . . . . .  
[31](#)  
 Authors' Addresses . . . . .  
[31](#)

**[1](#). Introduction**

This document defines two YANG [[RFC7950](#)] modules that augment the encapsulation choice YANG element defined in Interface Extensions YANG [[I-D.ietf-netmod-intf-ext-yang](#)] and the generic interfaces data model defined in [[RFC8343](#)]. The two modules provide configuration nodes to support classification of Ethernet/VLAN traffic to sub-interfaces, that can have interface based feature and service configuration applied to them.

The purpose of these models is to allow IETF defined forwarding protocols, such as IPv6 [[RFC2460](#)], Ethernet Pseudo Wires [[RFC4448](#)] and VPLS [[RFC4761](#)] [[RFC4762](#)] to be configurable via YANG when interoperating with VLAN tagged traffic received from an IEEE 802.1Q compliant bridge.

In the case of layer 2 Ethernet services, the flexible encapsulation module also supports flexible rewriting of the VLAN tags contained the in frame header.

For reference, a comparison between the sub-interface based YANG model documented in this draft and an IEEE 802.1Q bridge model is described in [Appendix A](#).

In summary, the YANG modules defined in this internet draft are:

if-l3-vlan.yang - Defines the model for basic classification of VLAN tagged traffic to L3 transport services

flexible-encapsulation.yang - Defines the model for flexible classification of Ethernet/VLAN traffic to L2 transport services

**[1.1](#). Terminology**

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in

[BCP](#)

[14 RFC 2119](#) [[RFC2119](#)] [RFC 8174](#) [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

Sub-interface: A sub-interface is a small augmentation of a regular interface in the standard YANG module for Interface Management that represents a subset of the traffic handled by its parent interface. As such, it supports both configuration and operational data using any other YANG models that augment or reference interfaces in the



normal way. It is defined in Interface Extensions YANG [[I-D.ietf-netmod-intf-ext-yang](#)].

## **1.2. Tree Diagrams**

Tree diagrams used in this document follow the notation defined in [[RFC8340](#)].

## **2. Objectives**

The primary aim of the YANG modules contained in this draft is to provide the core model that is required to implement VLAN transport services on router based devices that is fully compatible with IEEE 802.1Q compliant bridges.

A secondary aim is for the modules to be structured in such a way that they can be cleanly extended in future.

### **2.1. Interoperability with IEEE 802.1Q compliant bridges**

The modules defined in this document are designed to fully interoperate with IEEE 802.1Q compliant bridges. In particular, the models are restricted to only matching, adding, or rewriting the 802.1Q VLAN tags in frames in ways that are compatible with IEEE 802.1Q compliant bridges.

## **3. L3 Interface VLAN Model**

The L3 Interface VLAN model provides appropriate leaves for termination of an 802.1Q VLAN tagged segment to a sub-interface based

L3 service. It allows for termination of traffic with up to two 802.1Q VLAN tags.

The "if-l3-vlan" YANG module has the following structure:

```
module: ietf-if-l3-vlan
  augment /if:interfaces/if:interface/if-ext:encapsulation
    /if-ext:encaps-type:
  +--:(dot1q-vlan)
    +--rw dot1q-vlan
      +--rw outer-tag
        | +--rw tag-type      dot1q-tag-type
        | +--rw vlan-id      vlanid
      +--rw second-tag!
        +--rw tag-type      dot1q-tag-type
        +--rw vlan-id      vlanid
```





#### 4. Flexible Encapsulation Model

The Flexible Encapsulation model is designed to allow for the flexible provisioning of layer 2 services. It provides the capability to classify Ethernet/VLAN frames received on an Ethernet trunk interface to sub-interfaces based on the fields available in the layer 2 headers. Once classified to sub-interfaces, it provides the capability to selectively modify fields within the layer 2 headers before the frame is handed off to the appropriate forwarding code for further handling.

The model supports a common core set of layer 2 header matches based on the 802.1Q tag type and VLAN Ids contained within the header up to a tag stack depth of two tags.

The model supports flexible rewrites of the layer 2 frame header for data frames as they are processed on the interface. It defines a set

of standard tag manipulations that allow for the insertion, removal, or rewrite of one or two 802.1Q VLAN tags. The expectation is that manipulations are generally implemented in a symmetrical fashion, i.e. if a manipulation is performed on ingress traffic on an interface then the reverse manipulation is always performed on egress

traffic out of the same interface. However, the model also allows for asymmetrical rewrites, which may be required to implement some forwarding models (such as E-Tree).

The final aim for the model design is for it to be cleanly extensible

to add in additional match and rewrite criteria of the layer 2 header, such as matching on the source or destination MAC address, PCP or DEI fields in the 802.1Q tags, or the EtherType of the frame payload. Rewrites can also be extended to allow for modification of other fields within the layer 2 frame header.

The "flexible-encapsulation" YANG module has the following structure:

```
module: ietf-flexible-encapsulation
  augment /if:interfaces/if:interface/if-ext:encapsulation
    /if-ext:encaps-type:
      +--:(flexible)
        +--rw flexible
          +--rw match
            | +--rw (match-type)
            |   +--:(default)
            |   | +--rw default?           empty
            |   +--:(untagged)
            |   | +--rw untagged?         empty
```

| +--:(dot1q-priority-tagged)

Wilton, et al.  
5]

Expires May 7, 2020

[Page

```
| | +--rw dot1q-priority-tagged
| | | +--rw tag-type dot1q-types:dot1q-tag-type
| +--:(dot1q-vlan-tagged)
| | +--rw dot1q-vlan-tagged
| | | +--rw outer-tag
| | | | +--rw tag-type dot1q-tag-type
| | | | +--rw vlan-id union
| | | +--rw second-tag!
| | | | +--rw tag-type dot1q-tag-type
| | | | +--rw vlan-id union
| | | +--rw match-exact-tags? empty
+--rw rewrite {flexible-rewrites}?
| +--rw (direction)?
| | +--:(symmetrical)
| | | +--rw symmetrical
| | | | +--rw dot1q-tag-rewrite {dot1q-tag-rewrites}?
| | | | +--rw pop-tags? uint8
| | | | +--rw push-tags!
| | | | | +--rw outer-tag
| | | | | | +--rw tag-type dot1q-tag-type
| | | | | | +--rw vlan-id vlanid
| | | | | +--rw second-tag!
| | | | | +--rw tag-type dot1q-tag-type
| | | | | +--rw vlan-id vlanid
| | | +--:(asymmetrical) {asymmetric-rewrites}?
| | | +--rw ingress
| | | | +--rw dot1q-tag-rewrite {dot1q-tag-rewrites}?
| | | | +--rw pop-tags? uint8
| | | | +--rw push-tags!
| | | | | +--rw outer-tag
| | | | | | +--rw tag-type dot1q-tag-type
| | | | | | +--rw vlan-id vlanid
| | | | | +--rw second-tag!
| | | | | +--rw tag-type dot1q-tag-type
| | | | | +--rw vlan-id vlanid
| | | +--rw egress
| | | | +--rw dot1q-tag-rewrite {dot1q-tag-rewrites}?
| | | | +--rw pop-tags? uint8
| | | | +--rw push-tags!
| | | | | +--rw outer-tag
| | | | | | +--rw tag-type dot1q-tag-type
| | | | | | +--rw vlan-id vlanid
| | | | | +--rw second-tag!
| | | | | +--rw tag-type dot1q-tag-type
| | | | | +--rw vlan-id vlanid
+--rw local-traffic-default-encaps!
| +--rw outer-tag
| | +--rw tag-type dot1q-tag-type
```



```
|  +--rw vlan-id      vlanid
+--rw second-tag!
  +--rw tag-type      dot1q-tag-type
  +--rw vlan-id      vlanid
```

## **5. L3 Interface VLAN YANG Module**

This YANG module augments the encapsulation container defined in Interface Extensions YANG [[I-D.ietf-netmod-intf-ext-yang](#)].

```
<CODE BEGINS> file "ietf-if-l3-vlan@2019-11-04.yang"
module ietf-if-l3-vlan {
  yang-version 1.1;

  namespace "urn:ietf:params:xml:ns:yang:ietf-if-l3-vlan";

  prefix if-l3-vlan;

  import ietf-interfaces {
    prefix if;
  }

  import iana-if-type {
    prefix ianaift;
  }

  import ieee802-dot1q-types {
    prefix dot1q-types;
  }

  import ietf-if-extensions {
    prefix if-ext;
  }

  organization
    "IETF NETMOD (NETCONF Data Modeling Language) Working Group";

  contact
    "WG Web:   <http://tools.ietf.org/wg/netmod/>
    WG List:  <mailto:netmod@ietf.org>

    Editor:   Robert Wilton
              <mailto:rwilton@cisco.com>";

  description
    "This YANG module models L3 VLAN sub-interfaces
```



Copyright (c) 2019 IETF Trust and the persons identified as authors of the code. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, is permitted pursuant to, and subject to the license terms contained in, the Simplified BSD License set forth in [Section 4.c](#) of the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>).

This version of this YANG module is part of RFC XXXX (<https://www.rfc-editor.org/info/rfcXXXX>); see the RFC itself for full legal notices.";

```
revision 2019-11-04 {
  description "Latest draft revision";

  reference
    "Internet-Draft draft-ietf-netmod-sub-intf-vlan-model-06";
}

/*
 * Add support for the 802.1Q VLAN encapsulation syntax on layer 3
 * terminated VLAN sub-interfaces.
 */
augment "/if:interfaces/if:interface/if-ext:encapsulation/" +
  "if-ext:encaps-type" {
  when
    "derived-from-or-self(..if:type,
                          'ianaift:ethernetCsmacd') or
     derived-from-or-self(..if:type,
                          'ianaift:ieee8023adLag') or
     derived-from-or-self(..if:type,
                          'if-ext:ethSubInterface')" {
    description
      "Applies only to Ethernet-like interfaces and
       sub-interfaces";
  }

  description
    "Augment the generic interface encapsulation with an
     basic 802.1Q VLAN encapsulation for sub-interfaces.";

  /*
   * Matches a single VLAN Id, or a pair of VLAN Ids to classify
   * traffic into an L3 service.
   */
  case dot1q-vlan {
```





```
container dot1q-vlan {
  must
    'count(..../if-ext:forwarding-mode) = 0 or ' +
    'derived-from-or-self(..../if-ext:forwarding-mode,' +
    ' "if-ext:layer-3-forwarding")' {
    error-message
      "If the interface forwarding-mode leaf is set then it
      must be set to an identity that derives from
      layer-3-forwarding";

    description
      "The forwarding-mode leaf on an interface can
      optionally be used to enforce consistency of
      configuration";
  }
}
```

```
description
  "Match VLAN tagged frames with specific VLAN Ids";
container outer-tag {
  must
    'tag-type = "dot1q-types:s-vlan" or ' +
    'tag-type = "dot1q-types:c-vlan"' {

    error-message
      "Only C-VLAN and S-VLAN tags can be matched";

    description
      "For IEEE 802.1Q interoperability, only C-VLAN and
      S-VLAN tags can be matched";
  }
}
```

```
description
  "Classifies traffic using the outermost VLAN tag on the
  frame.";

uses dot1q-types:dot1q-tag-classifier-grouping;
}
```

```
container second-tag {
  must
    '../outer-tag/tag-type = "dot1q-types:s-vlan" and ' +
    'tag-type = "dot1q-types:c-vlan"' {

    error-message
      "When matching two tags, the outermost tag must be
      specified and of S-VLAN type and the second outermost
      tag must be of C-VLAN tag type";
  }
}
```



```
        description
            "For IEEE 802.1Q interoperability, when matching two
            tags, it is required that the outermost tag exists
and
            is an S-VLAN, and the second outermost tag is a
            C-VLAN";
        }

    presence "Also classify on the second outermost VLAN tag";

    description
        "Classifies traffic using the second outermost VLAN tag
        on the frame.";

    uses dot1q-types:dot1q-tag-classifier-grouping;
    }
}
}
}
}
}
}
}
}
}
}
<CODE ENDS>
```

## 6. Flexible Encapsulation YANG Module

This YANG module augments the encapsulation container defined in Interface Extensions YANG [[I-D.ietf-netmod-intf-ext-yang](#)].

This YANG module also augments the interface container defined in [[RFC8343](#)].

```
<CODE BEGINS> file "ietf-flexible-encapsulation@2019-11-04.yang"
module ietf-flexible-encapsulation {
    yang-version 1.1;
    namespace
        "urn:ietf:params:xml:ns:yang:ietf-flexible-encapsulation";

    prefix flex;

    import ietf-interfaces {
        prefix if;
    }

    import iana-if-type {
        prefix ianaift;
    }

    import ietf-if-extensions {
```



```
    prefix if-ext;
  }

import ieee802-dot1q-types {
  prefix dot1q-types;
}

organization
  "IETF NETMOD (NETCONF Data Modeling Language) Working Group";

contact
  "WG Web:   <http://tools.ietf.org/wg/netmod/>
  WG List:  <mailto:netmod@ietf.org>

  Editor:   Robert Wilton
            <mailto:rwilton@cisco.com>";

description
  "This YANG module describes interface configuration for flexible
  VLAN matches and rewrites.

  Copyright (c) 2019 IETF Trust and the persons identified as
  authors of the code. All rights reserved.

  Redistribution and use in source and binary forms, with or
  without modification, is permitted pursuant to, and subject to
  the license terms contained in, the Simplified BSD License set
  forth in Section 4.c of the IETF Trust's Legal Provisions
  Relating to IETF Documents
  (https://trustee.ietf.org/license-info).

  This version of this YANG module is part of RFC XXXX
  (https://www.rfc-editor.org/info/rfcXXXX); see the RFC itself
  for full legal notices.";

revision 2019-11-04 {
  description "Latest draft revision";

  reference
    "Internet-Draft draft-ietf-netmod-sub-intf-vlan-model-06";
}

feature flexible-rewrites {
  description
    "This feature indicates whether the network element supports
    specifying flexible rewrite operations";
}
```



```
feature asymmetric-rewrites {
  description
    "This feature indicates whether the network element supports
    specifying different rewrite operations for the ingress
    rewrite operation and egress rewrite operation.";
}

feature dot1q-tag-rewrites {
  description
    "This feature indicates whether the network element supports
    the flexible rewrite functionality specifying flexible 802.1Q
    tag rewrites";
}

/*
 * flexible-match grouping.
 *
 * This grouping represents a flexible match.
 *
 * The rules for a flexible match are:
 *   1. default, untagged, priority tag, or a stack of tags.
 *   - Each tag in the stack of tags matches:
 *     1. tag type (802.1Q or 802.1ad) +
 *     2. tag value:
 *       i. single tag
 *       ii. set of tag ranges/values.
 *       iii. "any" keyword
 */
grouping flexible-match {
  description "Flexible match";
  choice match-type {
    mandatory true;
    description "Provides a choice of how the frames may be
    matched";

    case default {
      description "Default match";
      leaf default {
        type empty;
        description
          "Default match. Matches all traffic not matched to any
          other peer sub-interface by a more specific
          encapsulation.";
      } // leaf default
    } // case default

    case untagged {
      description "Match untagged Ethernet frames only";
    }
  }
}
```





```
leaf untagged {
  type empty;
  description
    "Untagged match. Matches all untagged traffic.";
} // leaf untagged
} // case untagged

case dot1q-priority-tagged {
  description
    "Match 802.1Q priority tagged Ethernet frames only";

  container dot1q-priority-tagged {
    description "802.1Q priority tag match";
    leaf tag-type {
      type dot1q-types:dot1q-tag-type;
      mandatory true;
      description "The 802.1Q tag type of matched priority
        tagged packets";
    }
  }
}

case dot1q-vlan-tagged {
  container dot1q-vlan-tagged {
    description "Matches VLAN tagged frames";

    container outer-tag {
      must
        'tag-type = "dot1q-types:s-vlan" or ' +
        'tag-type = "dot1q-types:c-vlan"' {

        error-message
          "Only C-VLAN and S-VLAN tags can be matched";

        description
          "For IEEE 802.1Q interoperability, only C-VLAN and
            S-VLAN tags can be matched";
      }

      description
        "Classifies traffic using the outermost VLAN tag on
          the
            frame.";

      uses
        'dot1q-types:'+
        'dot1q-tag-ranges-or-any-classifier-grouping';
    }
  }
}
```



```
    container second-tag {
      must
        '../outer-tag/tag-type = "dot1q-types:s-vlan" and ' +
        'tag-type = "dot1q-types:c-vlan"' {

        error-message
          "When matching two tags, the outermost tag must be
            specified and of S-VLAN type and the second
            outermost tag must be of C-VLAN tag type";

        description
          "For IEEE 802.1Q interoperability, when matching two
            tags, it is required that the outermost tag exists
            and is an S-VLAN, and the second outermost tag is a
            C-VLAN";
      }

      presence "Also classify on the second VLAN tag";

      description
        "Classifies traffic using the second outermost VLAN
          on the frame.";

      uses
        'dot1q-types:' +
        'dot1q-tag-ranges-or-any-classifier-grouping';
    }

    leaf match-exact-tags {
      type empty;
      description
        "If set, indicates that all 802.1Q VLAN tags in the
          Ethernet frame header must be explicitly matched,
          the EtherType following the matched tags must not be
          a 802.1Q tag EtherType.  If unset then extra 802.1Q
          VLAN tags are allowed.";
    }
  }
} // encaps-type
}

/*
 * Grouping for tag-rewrite that can be expressed either
 * symmetrically, or in the ingress and/or egress directions
 * independently.
 */
grouping dot1q-tag-rewrite {
```



```
description "Flexible rewrite";
leaf pop-tags {
  type uint8 {
    range 1..2;
  }
  description "The number of tags to pop (or translate if used
in
                conjunction with push-tags)";
}

container push-tags {
  presence
    "802.1Q tags are pushed or translated";
  description "The 802.1Q tags to push (or translate if used in
                conjunction with pop-tags)";

  container outer-tag {
    must
      'tag-type = "dot1q-types:s-vlan" or ' +
      'tag-type = "dot1q-types:c-vlan"' {

      error-message
        "Only C-VLAN and S-VLAN tags can be pushed";

      description
        "For IEEE 802.1Q interoperability, only C-VLAN and S-
VLAN
        tags can be pushed";
    }

    description
      "The outermost VLAN tag to push onto the frame.";
    uses dot1q-types:dot1q-tag-classifier-grouping;
  }

  container second-tag {
    must
      '../outer-tag/tag-type = "dot1q-types:s-vlan" and ' +
      'tag-type = "dot1q-types:c-vlan"' {

      error-message
        "When pushing/rewriting two tags, the outermost tag must
outermost
        be specified and of S-VLAN type and the second
        tag must be of C-VLAN tag type";

      description
        "For IEEE 802.1Q interoperability, when pushing two
tags,
        it is required that the outermost tag exists and is an
        S-VLAN, and the second outermost tag is a C-VLAN";
    }
  }
}
```



```
        presence
            "In addition to the first tag, also push/rewrite a second
            VLAN tag.";

        description
            "The second outermost VLAN tag to push onto the frame.";

        uses dot1q-types:dot1q-tag-classifier-grouping;
    }
}

/*
 * Grouping for all flexible rewrites of fields in the L2 header.
 *
 * This currently only includes flexible tag rewrites, but is
 * designed to be extensible to cover rewrites of other fields in
 * the L2 header if required.
 */
grouping flexible-rewrite {
    description "Flexible rewrite";

    /*
     * Tag rewrite.
     *
     * All tag rewrites are formed using a combination of pop-tags
     * and push-tags operations.
     */
    container dot1q-tag-rewrite {
        if-feature dot1q-tag-rewrites;
        description "Tag rewrite. Translate operations are expressed
            as a combination of tag push and pop
operations.";
        uses dot1q-tag-rewrite;
    }
}
augment "/if:interfaces/if:interface/if-ext:encapsulation/" +
    "if-ext:encaps-type" {
    when
        "derived-from-or-self(..if:type,
            'ianaift:ethernetCsmacd') or
        derived-from-or-self(..if:type,
            'ianaift:ieee8023adLag') or
        derived-from-or-self(..if:type,
            'if-ext:ethSubInterface')" {
        description
            "Applies only to Ethernet-like interfaces and
            sub-interfaces";
    }
}
```





```
description
  "Add flexible match and rewrite for VLAN sub-interfaces";

/*
 * A flexible encapsulation allows for the matching of ranges
and
 * sets of VLAN Ids. The structure is also designed to be
 * extended to allow for matching/rewriting other fields within
 * the L2 frame header if required.
 */
case flexible {
  description "Flexible encapsulation and rewrite";
  container flexible {
    description "Flexible encapsulation and rewrite";

    container match {
      description
        "The match used to classify frames to this interface";
      uses flexible-match;
    }

    container rewrite {
      if-feature flexible-rewrites;
      description "L2 frame rewrite operations";
      choice direction {
        description
          "Whether the rewrite policy is symmetrical or
          asymmetrical";
        case symmetrical {
          container symmetrical {
            uses flexible-rewrite;
            description
              "Symmetrical rewrite. Expressed in the ingress
              direction, but the reverse operation is applied
to
              egress traffic";
          }
        }
      }
    }
  }
}

/*
 * Allow asymmetrical rewrites to be specified.
 */
case asymmetrical {
  if-feature asymmetric-rewrites;
  description "Asymmetrical rewrite";
  container ingress {
    uses flexible-rewrite;
    description "Ingress rewrite";
  }
  container egress {
```



```
        uses flexible-rewrite;
        description "Egress rewrite";
    }
}
}
}
}
/*
 * For encapsulations that match a range of VLANs (or Any),
 * allow configuration to specify the default 802.1Q VLAN
tag
 * values to use for any traffic that is locally sourced
from
 * an interface on the device.
 */
container local-traffic-default-encaps {
    presence
        "A local traffic default encapsulation has been
        specified";
    description
        "The 802.1Q VLAN tags to use by default for locally
        sourced traffic";

    container outer-tag {
        must
            'tag-type = "dot1q-types:s-vlan" or ' +
            'tag-type = "dot1q-types:c-vlan"' {

            error-message
                "Only C-VLAN and S-VLAN tags can be matched";

            description
                "For IEEE 802.1Q interoperability, only C-VLAN and
                S-VLAN tags can be matched";
        }

        description
            "The outermost VLAN tag for locally sourced traffic";

        uses dot1q-types:dot1q-tag-classifier-grouping;
    }

    container second-tag {
        must
            '../outer-tag/tag-type = "dot1q-types:s-vlan" and ' +
            'tag-type = "dot1q-types:c-vlan"' {

            error-message
                "When specifying two tags, the outermost tag must be
                specified and of S-VLAN type and the second outermost
```







```
<type>ianaift:ethernetCsmacd</type>
</interface>
<interface>
  <name>eth0.1</name>
  <type>ianaift:l2vlan</type>
  <if-cmn:parent-interface>eth0</if-cmn:parent-interface>
  <if-cmn:encapsulation>
    <dot1q-vlan
      xmlns="urn:ietf:params:xml:ns:yang:ietf-if-l3-vlan">
      <outer-tag>
        <tag-type>dot1q-types:s-vlan</tag-type>
        <vlan-id>10</vlan-id>
      </outer-tag>
      <second-tag>
        <tag-type>dot1q-types:c-vlan</tag-type>
        <vlan-id>20</vlan-id>
      </second-tag>
    </dot1q-vlan>
  </if-cmn:encapsulation>
  <ipv6 xmlns="urn:ietf:params:xml:ns:yang:ietf-ip">
    <forwarding>true</forwarding>
    <address>
      <ip>2001:db8::10</ip>
      <prefix-length>32</prefix-length>
    </address>
  </ipv6>
</interface>
<interface>
  <name>eth0.2</name>
  <type>ianaift:l2vlan</type>
  <if-cmn:parent-interface>eth0</if-cmn:parent-interface>
  <if-cmn:encapsulation>
    <dot1q-vlan
      xmlns="urn:ietf:params:xml:ns:yang:ietf-if-l3-vlan">
      <outer-tag>
        <tag-type>dot1q-types:s-vlan</tag-type>
        <vlan-id>11</vlan-id>
      </outer-tag>
    </dot1q-vlan>
  </if-cmn:encapsulation>
  <ipv6 xmlns="urn:ietf:params:xml:ns:yang:ietf-ip">
    <forwarding>true</forwarding>
    <address>
      <ip>2001:db8::11</ip>
      <prefix-length>32</prefix-length>
    </address>
  </ipv6>
</interface>
```





```
</interfaces>
</config>
```

## 7.2. Layer 2 sub-interfaces with L2VPN

This example illustrates a layer 2 sub-interface 'eth0.3' configured to match traffic with a S-VLAN tag of 10, and C-VLAN tag of 21; and both tags removed before the traffic is passed off to the L2VPN service.

It also illustrates another sub-interface 'eth1.0' under a separate physical interface configured to match traffic with a C-VLAN of 50, and the tag removed before traffic is given to any service. Sub-interface 'eth1.0' is not currently bound to any service and hence traffic classified to that sub-interface is dropped.

```
<?xml version="1.0" encoding="utf-8"?>
<config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <interfaces
    xmlns="urn:ietf:params:xml:ns:yang:ietf-interfaces"
    xmlns:ianaift="urn:ietf:params:xml:ns:yang:iana-if-type"
    xmlns:dot1q-types="urn:ieee:std:802.1Q:yang:ieee802-dot1q-types"
    xmlns:if-cmn="urn:ietf:params:xml:ns:yang:ietf-interfaces-common">
    <interface>
      <name>eth0</name>
      <type>ianaift:ethernetCsmacd</type>
    </interface>
    <interface>
      <name>eth0.3</name>
      <type>ianaift:l2vlan</type>
      <if-cmn:parent-interface>eth0</if-cmn:parent-interface>
      <if-cmn:encapsulation>
        <flexible
          xmlns="urn:ietf:params:xml:ns:yang:ietf-flexible-encapsulation">
          <match>
            <dot1q-vlan-tagged>
              <outer-tag>
                <tag-type>dot1q-types:s-vlan</tag-type>
                <vlan-id>10</vlan-id>
              </outer-tag>
              <second-tag>
                <tag-type>dot1q-types:c-vlan</tag-type>
                <vlan-id>21</vlan-id>
              </second-tag>
            </dot1q-vlan-tagged>
          </match>
```



```
        <rewrite>
          <symmetrical>
            <dot1q-tag-rewrite>
              <pop-tags>2</pop-tags>
            </dot1q-tag-rewrite>
          </symmetrical>
        </rewrite>
      </flexible>
    </if-cmn:encapsulation>
  </interface>
<interface>
  <name>eth1</name>
  <type>ianaift:ethernetCsmacd</type>
</interface>
<interface>
  <name>eth1.0</name>
  <type>ianaift:l2vlan</type>
  <if-cmn:parent-interface>eth0</if-cmn:parent-interface>
  <if-cmn:encapsulation>
    <flexible
xmlns="urn:ietf:params:xml:ns:yang:ietf-flexible-encapsulation">
      <match>
        <dot1q-vlan-tagged>
          <outer-tag>
            <tag-type>dot1q-types:c-vlan</tag-type>
            <vlan-id>50</vlan-id>
          </outer-tag>
        </dot1q-vlan-tagged>
      </match>
      <rewrite>
        <symmetrical>
          <dot1q-tag-rewrite>
            <pop-tags>1</pop-tags>
          </dot1q-tag-rewrite>
        </symmetrical>
      </rewrite>
    </flexible>
  </if-cmn:encapsulation>
</interface>
</interfaces>
<network-instances
xmlns="urn:ietf:params:xml:ns:yang:ietf-network-instance">
  <network-instance
xmlns:l2vpn="urn:ietf:params:xml:ns:yang:ietf-l2vpn">
    <name>p2p-l2-1</name>
    <description>Point to point L2 service</description>
    <l2vpn:type>l2vpn:vpws-instance-type</l2vpn:type>
    <l2vpn:signaling-type>
```



```
    l2vpn:ldp-signaling
  </l2vpn:signaling-type>
  <endpoint xmlns="urn:ietf:params:xml:ns:yang:ietf-l2vpn">
    <name>local</name>
    <ac>
      <name>eth0.3</name>
    </ac>
  </endpoint>
  <endpoint xmlns="urn:ietf:params:xml:ns:yang:ietf-l2vpn">
    <name>remote</name>
    <pw>
      <name>pw1</name>
    </pw>
  </endpoint>
  <vsi-root>
  </vsi-root>
</network-instance>
</network-instances>
<pseudowires
  xmlns="urn:ietf:params:xml:ns:yang:ietf-pseudowires">
  <pseudowire>
    <name>pw1</name>
    <configured-pw>
      <peer-ip>2001:db8::50</peer-ip>
      <pw-id>100</pw-id>
    </configured-pw>
  </pseudowire>
</pseudowires>
</config>
```

## **8. Acknowledgements**

The authors would particularly like to thank John Messenger, Glenn Parsons, and Dan Romascanu for their help progressing this draft.

The authors would also like to thank Alex Campbell, Eric Gray, Giles Heron, Marc Holness, Iftekhar Hussain, Neil Ketley, William Lupton, John Messenger, Glenn Parsons, Ludwig Pauwels, Joseph White, Vladimir Vassilev, and members of the IEEE 802.1 WG for their helpful reviews and feedback on this draft.

## **9. ChangeLog**

XXX, RFC Editor, please delete this change log before publication.



### **9.1. WG version -05**

- o Incorporate feedback from IEEE 802.1 WG, John Messenger in particular.
- o Adding must constraints to ensure outer tags are always matched to C-VLAN and S-VLAN tags.
- o Fixed bug where second tag could be matched without outer tag, and where tags must not be specified.

### **9.2. WG version -04**

- o Added examples

### **9.3. WG version -03**

- o Fix namespace bug in XPath identity references, removed extraneous 'dot1q-tag' containers.

### **9.4. WG version -02**

- o Use explicit containers for outer and inner tags rather than lists.

### **9.5. WG version -01**

- o Tweaked the abstract.
- o Removed unnecessary feature for the L3 sub-interface module.
- o Update the 802.1Qcp type references.
- o Remove extra tag container for L3 sub-interfaces YANG.

### **9.6. Version -04**

- o IEEE 802.1 specific types have been removed from the draft. These are now referenced from the 802.1Qcp draft YANG modules.
- o Fixed errors in the xpath expressions.

### **9.7. Version -03**

- o Incorporates feedback received from presenting to the IEEE 802.1 WG.





- o Updates the modules for double tag matches/rewrites to restrict the outer tag type to S-VLAN and inner tag type to C-VLAN.
- o Updates the introduction to indicate primary use case is for IETF forwarding protocols.
- o Updates the objectives to make IEEE 802.1Q bridge interoperability a key objective.

## **10. IANA Considerations**

This document defines two new YANG module and the authors politely request that IANA assigns unique names to the YANG module files contained within this draft, and also appropriate URIs in the "IETF XML Registry".

## **11. Security Considerations**

The YANG module defined in this memo is designed to be accessed via the NETCONF protocol [RFC 6241](#) [[RFC6241](#)]. The lowest NETCONF layer is the secure transport layer and the mandatory to implement secure transport is SSH [RFC 6242](#) [[RFC6242](#)]. The NETCONF access control model [RFC 6536](#) [[RFC6536](#)] provides the means to restrict access for particular NETCONF users to a pre-configured subset of all available NETCONF protocol operations and content.

There are a number of data nodes defined in this YANG module which are writable/creatable/deletable (i.e. config true, which is the default). These data nodes may be considered sensitive or vulnerable in some network environments. Write operations (e.g. edit-config) to these data nodes without proper protection can have a negative effect on network operations. These are the subtrees and data nodes and their sensitivity/vulnerability:

### **11.1. if-l3-vlan.yang**

The nodes in the if-l3-vlan YANG module are concerned with matching particular frames received on the network device to connect them to a layer 3 forwarding instance, and as such adding/modifying/deleting these nodes has a high risk of causing traffic to be lost because it is not being classified correctly, or is being classified to a separate sub-interface. The nodes, all under the subtree /interfaces/interface/encapsulation/dot1q-vlan, that are sensitive to this are:

- o outer-tag/tag-type

- o outer-tag/vlan-id

Wilton, et al.  
25]

Expires May 7, 2020

[Page

- o second-tag/tag-type
- o second-tag/vlan-id

### **11.2. flexible-encapsulation.yang**

There are many nodes in the flexible-encapsulation YANG module that are concerned with matching particular frames received on the network

device, and as such adding/modifying/deleting these nodes has a high risk of causing traffic to be lost because it is not being classified

correctly, or is being classified to a separate sub-interface. The nodes, all under the subtree /interfaces/interface/encapsulation/flexible/match, that are sensitive to this are:

- o default
- o untagged
- o dot1q-priority-tagged
- o dot1q-priority-tagged/tag-type
- o dot1q-vlan-tagged/outer-tag/vlan-type
- o dot1q-vlan-tagged/outer-tag/vlan-id
- o dot1q-vlan-tagged/second-tag/vlan-type
- o dot1q-vlan-tagged/second-tag/vlan-id

There are also many nodes in the flexible-encapsulation YANG module that are concerned with rewriting the fields in the L2 header for particular frames received on the network device, and as such adding/modifying/deleting these nodes has a high risk of causing traffic to be dropped or incorrectly processed on peer network devices, or it could cause layer 2 tunnels to go down due to a mismatch in negotiated MTU. The nodes, all under the subtree /interfaces/interface/encapsulation/flexible/rewrite, that are sensitive to this are:

- o symmetrical/dot1q-tag-rewrite/pop-tags
- o symmetrical/dot1q-tag-rewrite/push-tags/outer-tag/tag-type
- o symmetrical/dot1q-tag-rewrite/push-tags/outer-tag/vlan-id
- o symmetrical/dot1q-tag-rewrite/push-tags/second-tag/tag-type



- o symmetrical/dot1q-tag-rewrite/push-tags/second-tag/vlan-id
- o asymmetrical/ingress/dot1q-tag-rewrite/pop-tags
- o asymmetrical/ingress/dot1q-tag-rewrite/push-tags/outer-tag/tag-type
- o asymmetrical/ingress/dot1q-tag-rewrite/push-tags/outer-tag/vlan-id
- o asymmetrical/ingress/dot1q-tag-rewrite/push-tags/second-tag/tag-type
- o asymmetrical/ingress/dot1q-tag-rewrite/push-tags/second-tag/vlan-id
- o asymmetrical/egress/dot1q-tag-rewrite/pop-tags
- o asymmetrical/egress/dot1q-tag-rewrite/push-tags/outer-tag/tag-type
- o asymmetrical/egress/dot1q-tag-rewrite/push-tags/outer-tag/vlan-id
- o asymmetrical/egress/dot1q-tag-rewrite/push-tags/second-tag/tag-type
- o asymmetrical/egress/dot1q-tag-rewrite/push-tags/second-tag/vlan-id

Nodes in the flexible-encapsulation YANG module that are concerned with the VLAN tags to use for traffic sourced from the network element could cause protocol sessions (such as CFM) to fail if they are added, modified or deleted. The nodes, all under the subtree /interfaces/interface/flexible-encapsulation/local-traffic-default-encaps that are sensitive to this are:

- o outer-tag/vlan-type
- o outer-tag/vlan-id
- o second-tag/vlan-type
- o second-tag/vlan-id

## **12. References**

### **12.1. Normative References**



[I-D.ietf-netmod-intf-ext-yang]

Wilton, R., Ball, D., tsingh@juniper.net, t., and S. Sivaraj, "Common Interface Extension YANG Data Models", [draft-ietf-netmod-intf-ext-yang-07](#) (work in progress), March 2019.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC7224] Bjorklund, M., "IANA Interface Type YANG Module", [RFC 7224](#), DOI 10.17487/RFC7224, May 2014, <<https://www.rfc-editor.org/info/rfc7224>>.

[RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", [RFC 7950](#), DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[RFC8343] Bjorklund, M., "A YANG Data Model for Interface Management", [RFC 8343](#), DOI 10.17487/RFC8343, March 2018, <<https://www.rfc-editor.org/info/rfc8343>>.

## **12.2. Informative References**

[dot1Qcp] Holness, M., "IEEE 802.1Qcp-2018 Bridges and Bridged Networks - Amendment: YANG Data Model", 2018.

[I-D.ietf-bess-l2vpn-yang]

Shah, H., Brissette, P., Chen, I., Hussain, I., Wen, B., and K. Tiruveedhula, "YANG Data Model for MPLS-based L2VPN", [draft-ietf-bess-l2vpn-yang-10](#) (work in progress), July 2019.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), DOI 10.17487/RFC2460, December 1998, <<https://www.rfc-editor.org/info/rfc2460>>.

[RFC4448] Martini, L., Ed., Rosen, E., El-Aawar, N., and G. Heron, "Encapsulation Methods for Transport of Ethernet over MPLS Networks", [RFC 4448](#), DOI 10.17487/RFC4448, April 2006, <<https://www.rfc-editor.org/info/rfc4448>>.





- [RFC4761] Kompella, K., Ed. and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", [RFC 4761](#), DOI 10.17487/RFC4761, January 2007, <<https://www.rfc-editor.org/info/rfc4761>>.
- [RFC4762] Lasserre, M., Ed. and V. Kompella, Ed., "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", [RFC 4762](#), DOI 10.17487/RFC4762, January 2007, <<https://www.rfc-editor.org/info/rfc4762>>.
- [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, <<https://www.rfc-editor.org/info/rfc6241>>.
- [RFC6242] Wasserman, M., "Using the NETCONF Protocol over Secure Shell (SSH)", [RFC 6242](#), DOI 10.17487/RFC6242, June 2011, <<https://www.rfc-editor.org/info/rfc6242>>.
- [RFC6536] Bierman, A. and M. Bjorklund, "Network Configuration Protocol (NETCONF) Access Control Model", [RFC 6536](#), DOI 10.17487/RFC6536, March 2012, <<https://www.rfc-editor.org/info/rfc6536>>.
- [RFC8340] Bjorklund, M. and L. Berger, Ed., "YANG Tree Diagrams", [BCP 215](#), [RFC 8340](#), DOI 10.17487/RFC8340, March 2018, <<https://www.rfc-editor.org/info/rfc8340>>.

## [Appendix A](#). Comparison with the IEEE 802.1Q Configuration Model

In addition to the sub-interface based YANG model proposed here, the IEEE 802.1Q working group has developed a YANG model for the configuration of 802.1Q VLANs. This raises the valid question as to whether the models overlap and whether it is necessary or beneficial to have two different models for superficially similar constructs. This section aims to answer that question by summarizing and comparing the two models.

### [A.1](#). Sub-interface based configuration model overview

The key features of the sub-interface based configuration model can be summarized as:

- o The model is primarily designed to enable layer 2 and layer 3 services on Ethernet interfaces that can be defined in a very flexible way to meet the varied requirements of service providers.



- o Traffic is classified from an Ethernet-like interface to sub-interfaces based on fields in the layer 2 header. This is often based on VLAN Ids contained in the frame, but the model is extensible to other arbitrary fields in the frame header.
- o Sub-interfaces are just a type of if:interface and hence support any feature configuration YANG models that can be applied generally to interfaces. For example, QoS or ACL models that reference if:interface can be applied to the sub-interfaces, or the sub-interface can be used as an Access Circuit in L2VPN or L3VPN models that reference if:interface.
- o In the sub-interface based configuration model, the classification of traffic arriving on an interface to a given sub-interface, based on fields in the layer 2 header, is completely independent of how the traffic is forwarded. The sub-interface can be referenced (via references to if:interface) by other models that specify how traffic is forwarded; thus sub-interfaces can support multiple different forwarding paradigms, including but not limited to: layer 3 (IPv4/IPv6), layer 2 pseudowires (over MPLS or IP), VPLS instances, EVPN instance.
- o The model is flexible in the scope of the VLAN Identifier space. I.e. by default VLAN Ids can be scoped locally to a single Ethernet-like trunk interface, but the scope is determined by the forwarding paradigm that is used.

## **A.2. IEEE 802.1Q Bridge Configuration Model Overview**

The key features of the IEEE 802.1Q bridge configuration model can be summarized as:

- o Each VLAN bridge component has a set of Ethernet interfaces that are members of that bridge. Sub-interfaces are not used, nor required in the 802.1Q bridge model.
- o Within a VLAN bridge component, the VLAN tag in the packet is used, along with the destination MAC address, to determine how to forward the packet. Other forwarding paradigms are not supported by the 802.1Q model.
- o Classification of traffic to a VLAN bridge component is based only on the Ethernet interface that it arrived on.
- o VLAN Identifiers are scoped to a VLAN bridge component. Often devices only support a single bridge component and hence VLANs are scoped globally within the device.



- o Feature configuration is specified in the context of the bridge, or particular VLANs on a bridge.

### **A.3. Possible Overlap Between the Two Models**

Both models can be used for configuring similar basic layer 2 forwarding topologies. The 802.1Q bridge configuration model is optimised for configuring Virtual LANs that span across enterprises and data centers.

The sub-interface model can also be used for configuring equivalent Virtual LAN networks that span across enterprises and data centers, but often requires more configuration to be able to configure the equivalent constructs to the 802.1Q bridge model.

The sub-interface model really excels when implementing flexible L2 and L3 services, where those services may be handled on the same physical interface, and where the VLAN Identifier is being solely used to identify the customer or service that is being provided rather than a Virtual LAN. The sub-interface model provides more flexibility as to how traffic can be classified, how features can be applied to traffic streams, and how the traffic is to be forwarded.

Conversely, the 802.1Q bridge model can also be use to implement L2 services in some scenarios, but only if the forwarding paradigm being used to implement the service is the native Ethernet forwarding specified in 802.1Q - other forwarding paradigms such as pseudowires or VPLS are not supported. The 802.1Q bridge model does not implement L3 services at all, although this can be partly mitigated by using a virtual L3 interface construct that is a separate logical Ethernet-like interface which is a member of the bridge.

In conclusion, it is valid for both of these models to exist since they have different deployment scenarios for which they are optimized. Devices may choose which of the models (or both) to implement depending on what functionality the device is being designed for.

#### Authors' Addresses

Robert Wilton (editor)  
Cisco Systems

Email: [rwilton@cisco.com](mailto:rwilton@cisco.com)



Internet-Draft  
2019

Sub-interface VLAN YANG

November

David Ball  
Cisco Systems

Email: [daviball@cisco.com](mailto:daviball@cisco.com)

Tapraj Singh  
Cisco Systems

Email: [tapsingh@cisco.com](mailto:tapsingh@cisco.com)

Selvakumar Sivaraj  
Juniper Networks

Email: [ssivaraj@juniper.net](mailto:ssivaraj@juniper.net)

