

Internet Working Group
Internet-Draft
Intended status: Standards Track

Y. Jiang
X. He
Huawei
W. Cheng
China Mobile
March 1, 2019

Expires: September 2019

YANG Data Model for FlexE Interface Management
draft-jiang-ccamp-flex-e-yang-00

Abstract

This document defines a YANG data model for the configuration of FlexE 2.0 interface, and its FlexE clients. The YANG module in this document conforms to the Network Management Datastore Architecture (NMDA).

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 September 1, 2019.

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 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in

Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
1.1.	Conventions used in this document	3
1.2.	Terminology	3
2.	YANG model hierarchy for FlexE interface	4
3.	YANG Module for FlexE interface	6
4.	Security Considerations	12
5.	IANA Considerations	13
6.	References	13
6.1.	Normative References	13
6.2.	Informative References	14
7.	Acknowledgments	14

1. Introduction

The Flex Ethernet (FlexE) 2.0 Implementation Agreement [FLEXE] defined by the OIF provides the support of a variety of Ethernet MAC rates that may or may not correspond to any existing Ethernet PHY rate. This includes MAC rates that are both greater than (through bonding) and less than (through sub-rate and channelization) the Ethernet PHY rates used to carry FlexE. FlexE 2.0 further supports the bonding of 200GBASE-R PHYs or the bonding of 400GBASE-R PHYs.

In the FlexE, multiple Ethernet PHYs (each PHY can further consist of one or more FlexE Instances) are bonded into a FlexE Group, and the total capacity of the FlexE Group is represented as a collection of slots (e.g., each slot has a granularity of 5Gbps or 25Gbps). Based on their bandwidth needs, FlexE Clients are each mapped into one or more slots. The FlexE mechanism operates using a calendar consisting of these slots.

This calendar is partitioned into sub-calendars for each FlexE instance. For example, the calendar for a FlexE Group composed of n 100G PHYs is partitioned into 20n slots (each representing 5Gbps of bandwidth when the slot granularity is 5Gbps).

This document defines a YANG data model for the configuration of a Flex Ethernet interface (i.e., FlexE Group). The data model is augmented based on the generic interfaces data model as defined in [RFC8343], the FlexE attributes are based on the FlexE 2.0

Implementation Agreement as specified in [FLEXE]. With the help of this YANG module, the FlexE Groups can be managed just as network interfaces on a network device (e.g., a router or bridge).

The YANG module in this document conforms to the Network Management Datastore Architecture (NMDA) [RFC8342].

1.1. Conventions used in this document

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 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1.2. Terminology

Most terminologies used in this document are extracted from [FLEXE].

FlexE: Flex Ethernet

FlexE Client: An Ethernet flow based on a MAC data rate that may or may not correspond to any Ethernet PHY rate.

FlexE Group: A FlexE Group is composed of from 1 to n Ethernet PHYs.

FlexE Instance: A FlexE Instance is a unit of information consisting of 100G of capacity able to carry FlexE Client data, together with its associated overhead.

Ethernet PHY: an entity representing Ethernet Physical Coding Sublayer (PCS), Physical Media Attachment (PMA), and Physical Media Dependent (PMD) layers. Each PHY is consisted of one or more FlexE Instance (e.g., a 400GBASE-R PHY has four FlexE Instances)

FlexE Calendar: The total capacity of a FlexE Group is represented as a collection of slots. The calendar for a FlexE Group composed of n 100G PHYs is represented as an array of 20n slots (each representing 5Gbps of bandwidth). This calendar is partitioned into sub-calendars, e.g., a sub-calendar with 20 slots per 100G PHY.

2. YANG model hierarchy for FlexE interface

This section describes the hierarchy of the YANG module for FlexE interface management.

Configuration and status of FlexE interface information include:

- flexe-group specifies a FlexE group consisting of multiple PHYs.
- flexe-client-list specifies a list of FlexE client, each client is mapped to some slots in this FlexE group. flexe-client-status indicates whether there is any fault in any mapped slot for this client.

The readers are assumed to be familiar with FlexE 2.0, as all FlexE terminologies are described in details in [FLEXE].

In order to simplify the YANG module of the FlexE interface and to follow the YANG style of terminology, neither sub-calendar nor calendar in FlexE are modelled explicitly. However, a sub-calendar is represented by a calendar-slot-list per instance, and calendar is actually a conglomerate of all the slots in calendar slot lists for all FlexE instances in all FlexE PHYs of this FlexE Group.

A simplified YANG tree diagram [RFC8340] representing the data model is typically used by YANG modules. This document uses the same tree diagram syntax as described in [RFC8340].

```

module: ietf-flexe
  augment /if:interfaces/if:interface:
    +--rw flexe-group
      +--rw group-number?          uint32
      +--rw slot-granularity?      slot-granularity-enumeration
      +--rw flexe-phy-type?        flexe-phy-enumeration
      +--rw flexe-calendar-inuse?  calendar-enumeration
      +--rw flexe-phy-list* [phy-number]
        +--rw phy-number          uint8
        +--rw flexe-inst-list* [instance-number]
          +--rw instance-number    uint8
          +--rw calendar-slot-list* [slot-id]
            +--rw slot-id          uint8
            +--rw flexe-slot-status? slot-status-enumeration
      +--rw flexe-client-list* [client-id]
        +--rw client-id            uint16
        +--rw group-number?        uint32

```

```

+--rw mapped-slot-list* [mapped-slot-id]
|   +--rw mapped-slot-id      uint8
|   +--rw mapped-phy-number?   uint8
|   +--rw mapped-inst-number?  uint8
+--ro flexe-client-status?     uint8

```

3. YANG Module for FlexE interface

This module imports typedef "interface-ref" from [RFC8343].

```
<CODE BEGINS> file "ietf-flexex@2019-02-25.yang"
module ietf-flexex {
  yang-version 1.1;
  namespace "urn:ietf:params:xml:ns:yang:ietf-flexex";
  prefix "flexex";

  import iana-if-type {
    prefix ianaift;
  }
  import ietf-interfaces {
    prefix if;
    reference
      "RFC8343: A YANG Data Model for Interface Management";
  }

  organization "IETF CCAMP Working Group";
  contact
    "WG Web:  http://tools.ietf.org/wg/ccamp/
    WG List:  <mailto:ccamp@ietf.org>
    Author:   Yuanlong Jiang
              <mailto:jiangyuanlong@huawei.com>
    Author:   Xiang He
              <mailto:hexiang@huawei.com>
    Author:   Weiqiang Cheng
              <mailto:chengweiqiang@chinamobile.com>";

  description
    "This YANG module defines a data model for the configuration
    of FlexE interface.";

  revision "2019-02-25" {
    description "Initial version";
    reference
      "draft-jiang-ccamp-flexex-yang-00: YANG Data Model for FlexE
      Interface Management ";
  }

  typedef slot-granularity-enumeration {
    type enumeration {
      enum slot-5g {
        value 1;
        description
          "5Gbps per slot.";
      }
    }
  }
```

```
    enum slot-25g {
        value 2;
        description
            "25Gbps per slot.";
    }
    enum slot-others {
        value 254;
        description
            "Other type of granularities per slot.";
    }
}
description
    "The bandwidth granularity of a slot. Options for this
    enumeration are specified by the OIF standard, currently only
    5G and 25G are defined.";
reference
    "OIF Flex 2.0: Section 6.7";
}

typedef slot-status-enumeration {
    type enumeration {
        enum unavailable {
            value 1;
            description
                "slot is unavailable for FlexE client.";
        }
        enum unused {
            value 2;
            description
                "slot is unused.";
        }
        enum used {
            value 3;
            description
                "slot is used.";
        }
    }
}
description
    "The status of a slot. Options for this enumeration are
    specified by the OIF standard, 'used' is implicit.";
reference
    "OIF Flex 2.0: Section 7.3.4 and Section 7.6";
}

typedef flexe-phy-enumeration {
    type enumeration {
        enum 'PHY-100GBASE-R' {
```

```
        value 1;
        description
            "100GBASE-R PHY, as defined in FlexE 1.0.";
    }
    enum 'PHY-200GBASE-R' {
        value 2;
        description
            "200GBASE-R PHY, as defined in FlexE 2.0.";
    }
    enum 'PHY-400GBASE-R' {
        value 3;
        description
            "400GBASE-R PHY, as defined in FlexE 2.0.";
    }
}

description
    "The current type of PHYs bonded in a FlexE Group. Values for
    this enumeration are specified by the OIF standard.";
reference
    "OIF Flex 2.0: Section 5.2.1.5";
}

typedef calendar-enumeration {
    type enumeration {
        enum "CALENDAR-A" {
            description "Using Calendar A";
        }
        enum "CALENDAR-B" {
            description "Using Calendar B";
        }
    }
}
description
    "FlexE Calendar in use, values for this enumeration
    are specified by the OIF standard.";
reference
    "OIF Flex 2.0: Section 6.7";
}

augment "/if:interfaces/if:interface" {

    when "if:type = 'ianaift:flexethernet'" {
        description "Applies to Flex Ethernet interfaces";
    }

    description
```


"Augment interface model with OIF Flex Ethernet interface specific configuration nodes. Each FlexE interface represents a FlexE Group configured in a device.";

```
container flexe-group {
  description
    "The struct containing all FlexE related configuration
    (see OIF FlexE 2.0 Section 6.1).
    Note that max number of FlexE groups in a network is 63.";

  leaf group-number {
    type uint32 {
      range "1..1048574";
    }
    description
      "FlexE Group Number as defined in Section 7.3.6 of FlexE
      2.0.";
  }

  leaf slot-granularity {
    type slot-granularity-enum;
    default "slot-5g";
    description
      "The granularity of a slot in a FlexE group.";
  }

  leaf flexe-phy-type {
    type flexe-phy-enum;
    default "PHY-100GBASE-R";
    description
      "The type of PHYs bonded in a FlexE Group.";
  }

  leaf flexe-calendar-inuse {
    type calendar-enum;
    default "CALENDAR-A";
    description
      "The FlexE Calendar in use for a FlexE Group.";
  }

  list flexe-phy-list {
    key "phy-number";
    description
      "List of PHYs bonded in a FlexE group per FlexE 2.0.";
  }
}
```

```
leaf phy-number {
  type uint8 {
    range "1 .. 254";
  }
  description
    "PHY number of a FlexE PHY.
    If PHY type is 100GBASE-R, phy-number is [1,254].
    If PHY type is 200GBASE-R, phy-number is [1,126].
    If PHY type is 400GBASE-R, phy-number is [1, 62].";
}

list flexe-inst-list {
  key "instance-number";
  description
    "List of instances in a FlexE PHY as defined in OIF.
    Max elements of flexe-inst-list in a FlexE PHY:
    If PHY type is 100GBASE-R, max-elements is 1.
    If PHY type is 200GBASE-R, max-elements is 2.
    If PHY type is 400GBASE-R, max-elements is 4.";

  leaf instance-number {
    type uint8;
    description
      "Instance number of an instance. Its range relies on
      phy-number of the PHY wherein this instance belongs:
      If PHY type is 100GBASE-R:
        instance-number=phy-number
      If PHY type is 200GBASE-R:
        instance-number=2*phy-number+[0, 1]
      If PHY type is 400GBASE-R:
        instance-number=4*phy-number+[0, 3] ";
  }

  list calendar-slot-list {
    key "slot-id";
    leaf slot-id {
      type uint8;
      description
        "slot id of a slot in an instance.";
    }
    description
      "List of slots in a FlexE instance, i.e., sub-
      calendar per instance in FlexE 2.0. Max elements of
      calendar-slot-list for a FlexE instance is:
      If slot-granularity=slot-5g, max-elements is 20.
      If slot-granularity=slot-25g, max-elements is 4.";
  }
}
```

```
        leaf flexe-slot-status {
            type slot-status-enumeration;
            default unused;
            description
                "Slot status of a slot in an instance.";
        }
    } //calendar-slot-list
} //flexe-inst-list
} //flexe-phy-list

list flexe-client-list {

    key "client-id";
    description
        "List of FlexE clients in a FlexE Group.";
    leaf client-id {
        type uint16;
        description
            "FlexE client ID as defined in FlexE IA.";
    }
    leaf group-number {
        type uint32 {
            range "1..1048574";
        }
        description
            "FlexE Group Number of the FlexE group for this client.";
    }

    list mapped-slot-list {
        key "mapped-slot-id";
        description
            "List of mapped-slots for a FlexE client.";
        leaf mapped-slot-id {
            type uint8;
            description
                "Slot id of a slot in an instance for a client.";
        }
        leaf mapped-phy-number {
            type uint8;
            description
                "PHY number of a slot for a client.";
        }
        leaf mapped-inst-number {
            type uint8;
            description
                "Instance number of a slot for a client.";
        }
    }
}
```

```
    } // mapped-slot-list

    leaf flexe-client-status {
        type uint8;
        config false;
        description
            "Fault status for a client indicated in its mapped
            slots. If any slot is in fault, the client status
            is indicated in fault. Status includes:
            OK, Local Fault, Remote Fault and etc.";
    }

    } //flexe-client-list
} //flexe-group
} //augment
}

<CODE ENDS>
```

4. Security Considerations

The YANG module specified in this document defines a schema for data that is designed to be accessed via network management protocols such as NETCONF [RFC6241] or RESTCONF [RFC8040]. The lowest NETCONF layer is the secure transport layer, and the mandatory-to-implement secure transport is Secure Shell (SSH) [RFC6242]. The lowest RESTCONF layer is HTTPS, and the mandatory-to-implement secure transport is TLS [RFC8446].

The NETCONF access control model [RFC8341] provides the means to restrict access for particular NETCONF or RESTCONF users to a preconfigured subset of all available NETCONF or RESTCONF protocol operations and content.

There are a number of data nodes defined in this YANG module are writable, and the involved subtrees that are sensitive include:

/flexe/flexe-group-list specifies a list of FlexE Group, and each group consists of multiple PHYs.

/flexe/flexe-client-list specifies a list of FlexE Client, and each client is mapped to some slots in a FlexE Group.

Write operations (e.g., edit-config) to these data nodes without proper protection can have a negative effect on network operations. Specifically, an inappropriate configuration of them may cause an

interrupt of a client flow or even break down of a whole FlexE interface.

5. IANA Considerations

It is proposed that IANA register the following URI in the "IETF XML registry" [RFC3688]:

URI: urn:ietf:params:xml:ns:yang:ietf-flex

Registrant Contact: The IESG

XML: N/A; the requested URI is an XML namespace

It is proposed that IANA register the following YANG module in the "YANG Module Names" registry:

Name: ietf-flex

Namespace: urn:ietf:params:xml:ns:yang:ietf-flex

Prefix: flex

Reference: this document

It is proposed that IANA register a new IANAifType TBD for the interface type of Flex Ethernet in the "IANA Interface Type YANG Module" [RFC7224].

6. References

6.1. Normative References

- [FLEXE] OIF, "Flex Ethernet 2.0 Implementation Agreement", FlexE 2.0, June 2018
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997
- [RFC3688] Mealling, M., "The IETF XML Registry", RFC 3688, January 2004
- [RFC6241] Enns, R., Bjorklund, M., Schoenwaelder, J., and Bierman, A., "Network Configuration Protocol (NETCONF)", RFC 6241, June 2011
- [RFC6242] Wasserman, M., "Using the NETCONF Protocol over Secure Shell (SSH)", RFC 6242, June 2011
- [RFC7224] Bjorklund, M., "IANA Interface Type YANG Module", RFC 7224, May 2014

- [RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, January 2017
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, May 2017
- [RFC8341] Bierman, A. and Bjorklund, M., "Network Configuration Protocol (NETCONF) Access Control Model", RFC 8341, March 2018
- [RFC8342] Bjorklund, M., Schoenwaelder, J., Shafer, P., Watsen, K., and R. Wilton, "Network Management Datastore Architecture (NMDA)", RFC 8342, March 2018
- [RFC8343] Bjorklund, M., "A YANG Data Model for Interface Management", RFC 8343, March 2018
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, August 2018

6.2. Informative References

- [RFC8340] Bjorklund, M., and Berger, L., "YANG Tree Diagrams", RFC 8340, March 2018

7. Acknowledgments

The authors would like to thank Zitao Wang and Zhongjian Zhang for discussions on YANG.

Authors' Addresses

Yuanlong Jiang
Huawei Technologies Co., Ltd.
Bantian, Longgang district
Shenzhen 518129, China
Email: jiangyuanlong@huawei.com

Xiang He
Huawei Technologies Co., Ltd.
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
Email: hexiang9@huawei.com

Weiqiang Cheng
China Mobile
No.32 Xuanwumen West Street
Beijing 100053, China
Email: chengweiqiang@chinamobile.com

