

NETMOD WG  
Internet-Draft  
Intended status: Standards Track  
Expires: December 27, 2015

D. Bogdanovic  
Juniper Networks  
K. Sreenivasa  
Brocade Communications System  
L. Huang  
Juniper Networks  
D. Blair  
Cisco Systems  
June 25, 2015

**Network Access Control List (ACL) YANG Data Model**  
**draft-ietf-netmod-acl-model-03**

**Abstract**

This document describes a data model of Access Control List (ACL) basic building blocks.

**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 December 27, 2015.

**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

|                             |   |                    |
|-----------------------------|---|--------------------|
| <a href="#">1.</a>          | <a href="#">Introduction . . . . .</a>                                    | <a href="#">2</a>  |
| <a href="#">1.1.</a>        | <a href="#">Definitions and Acronyms . . . . .</a>                        | <a href="#">3</a>  |
| <a href="#">2.</a>          | <a href="#">Problem Statement . . . . .</a>                               | <a href="#">3</a>  |
| <a href="#">3.</a>          | <a href="#">Design of the ACL Model . . . . .</a>                         | <a href="#">4</a>  |
| <a href="#">3.1.</a>        | <a href="#">ACL Modules . . . . .</a>                                     | <a href="#">4</a>  |
| <a href="#">4.</a>          | <a href="#">ACL YANG Models . . . . .</a>                                 | <a href="#">6</a>  |
| <a href="#">4.1.</a>        | <a href="#">IETF Access Control List module . . . . .</a>                 | <a href="#">6</a>  |
| <a href="#">4.2.</a>        | <a href="#">IETF-PACKET-FIELDS module . . . . .</a>                       | <a href="#">10</a> |
| <a href="#">4.3.</a>        | <a href="#">An ACL Example . . . . .</a>                                  | <a href="#">15</a> |
| <a href="#">4.4.</a>        | <a href="#">Port Range Usage Example . . . . .</a>                        | <a href="#">16</a> |
| <a href="#">5.</a>          | <a href="#">Linux nftables . . . . .</a>                                  | <a href="#">17</a> |
| <a href="#">6.</a>          | <a href="#">Security Considerations . . . . .</a>                         | <a href="#">18</a> |
| <a href="#">7.</a>          | <a href="#">IANA Considerations . . . . .</a>                             | <a href="#">18</a> |
| <a href="#">8.</a>          | <a href="#">Acknowledgements . . . . .</a>                                | <a href="#">19</a> |
| <a href="#">9.</a>          | <a href="#">References . . . . .</a>                                      | <a href="#">19</a> |
| <a href="#">9.1.</a>        | <a href="#">Normative References . . . . .</a>                            | <a href="#">19</a> |
| <a href="#">9.2.</a>        | <a href="#">Informative References . . . . .</a>                          | <a href="#">19</a> |
| <a href="#">Appendix A.</a> | <a href="#">Extending ACL model examples . . . . .</a>                    | <a href="#">20</a> |
| <a href="#">A.1.</a>        | <a href="#">Example of extending existing model for route filtering .</a> | <a href="#">20</a> |
| <a href="#">A.2.</a>        | <a href="#">A company proprietary module example . . . . .</a>            | <a href="#">22</a> |
| <a href="#">A.3.</a>        | <a href="#">Attaching Access Control List to interfaces . . . . .</a>     | <a href="#">24</a> |
|                             | <a href="#">Authors' Addresses . . . . .</a>                              | <a href="#">26</a> |

## [1. Introduction](#)

Access Control List (ACL) is one of the basic elements to configure device forwarding behavior. It is used in many networking concepts such as Policy Based Routing, Firewalls etc.

An ACL is an ordered set of rules that is used to filter traffic on a networking device. Each rule is represented by an Access Control Entry (ACE).

Each ACE has a group of match criteria and a group of action criteria.

The match criteria consist of a tuple of packet header match criteria and metadata match criteria.

- o Packet header matches apply to fields visible in the packet such as address or class of service or port numbers.



- o Metadata matches apply to fields associated with the packet but not in the packet header such as input interface or overall packet length

The actions specify what to do with the packet when the matching criteria is met. These actions are any operations that would apply to the packet, such as counting, policing, or simply forwarding. The list of potential actions is endless depending on the innovations of the networked devices.

Access Control List is also widely known as ACL (pronounce as [ak-uh l]) or Access List. In this document, Access Control List, ACL and Access List are interchangeable.

### **1.1. Definitions and Acronyms**

ACE: Access Control Entry

ACL: Access Control List

AFI: Address Field Identifier

DSCP: Differentiated Services Code Point

ICMP: Internet Control Message Protocol

IP: Internet Protocol

IPv4: Internet Protocol version 4

IPv6: Internet Protocol version 6

MAC: Media Access Control

TCP: Transmission Control Protocol

## **2. Problem Statement**

This document defines a YANG [[RFC6020](#)] data model for the configuration of ACLs. It is very important that model can be easily reused between vendors and between applications.

ACL implementations in every device may vary greatly in terms of the filter constructs and actions that they support. Therefore this draft proposes a simple model that can be augmented by standard extensions and vendor proprietary models.



### **3. Design of the ACL Model**

Although different vendors have different ACL data models, there is a common understanding of what access control list (ACL) is. A network system usually have a list of ACLs, and each ACL contains an ordered list of rules, also known as access list entries - ACEs. Each ACE has a group of match criteria and a group of action criteria. The match criteria consist of packet header matching and metadata matching. Packet header matching applies to fields visible in the packet such as address or class of service or port numbers. Metadata matching applies to fields associated with the packet, but not in the packet header such as input interface, packet length, or source or destination prefix length. The actions can be any sort of operation from logging to rate limiting or dropping to simply forwarding. Actions on the first matching ACE are applied with no processing of subsequent ACEs. The model also includes a container to hold overall operational state for each ACL and operational state for each ACE. One ACL can be applied to multiple targets within the device, such as interfaces of a networked device, applications or features running in the device, etc. When applied to interfaces of a networked device, the ACL is applied in a direction which indicates if it should be applied to packet entering (input) or leaving the device (output). An example in the appendix shows how to express it in YNAG model.

This draft tries to address the commonalities between all vendors and create a common model, which can be augmented with proprietary models. The base model is very simple and with this design we hope to achieve needed flexibility for each vendor to extend the base model.

#### **3.1. ACL Modules**

There are two YANG modules in the model. The first module, "ietf-access-control-list", defines generic ACL aspects which are common to all ACLs regardless of their type or vendor. In effect, the module can be viewed as providing a generic ACL "superclass". It imports the second module, "ietf-packet-fields". The match container in "ietf-access-control-list" uses groupings in "ietf-packet-fields". If there is a need to define new "matches" choice, such as IPFIX [[RFC5101](#)], the container "matches" can be augmented.



```

module: ietf-access-control-list
+--rw access-lists
+--rw acl* [acl-name]
  +--ro acl-oper-data
  +--rw access-list-entries
    | +--rw ace* [rule-name]
    |   +--rw matches
    |     | +--rw (ace-type)?
    |     | | +--:(ace-ip)
    |     | | | +--rw (ace-ip-version)?
    |     | | | | +--:(ace-ipv4)
    |     | | | | +--rw destination-ipv4-network?  inet:ipv4-prefix
    |     | | | | +--rw source-ipv4-network?      inet:ipv4-prefix
    |     | | | | +--:(ace-ipv6)
    |     | | | | +--rw destination-ipv6-network?  inet:ipv6-prefix
    |     | | | | +--rw source-ipv6-network?      inet:ipv6-prefix
    |     | | | | +--rw flow-label?                inet:ipv6-flow-label
    |     | | | | +--rw dscp?                      inet:dscp
    |     | | | | +--rw protocol?                  uint8
    |     | | | | +--rw source-port-range
    |     | | | | | +--rw lower-port?  inet:port-number
    |     | | | | | +--rw upper-port?  inet:port-number
    |     | | | | +--rw destination-port-range
    |     | | | | | +--rw lower-port?  inet:port-number
    |     | | | | | +--rw upper-port?  inet:port-number
    |     | | | +--:(ace-eth)
    |     | | |   +--rw destination-mac-address?    yang:mac-address
    |     | | |   +--rw destination-mac-address-mask? yang:mac-address
    |     | | |   +--rw source-mac-address?        yang:mac-address
    |     | | |   +--rw source-mac-address-mask?   yang:mac-address
    |     | +--rw input-interface?                  string
    |     | +--rw absolute-time
    |     |   +--rw start?    yang:date-and-time
    |     |   +--rw end?      yang:date-and-time
    |     |   +--rw active?   boolean
    |     +--rw actions
    |       | +--rw (packet-handling)?
    |       |   +--:(deny)
    |       |   | +--rw deny?    empty
    |       |   +--:(permit)
    |       |   | +--rw permit?  empty
    |       +--ro ace-oper-data
    |       | +--ro match-counter?  yang:counter64
    |       +--rw rule-name        string
  +--rw acl-name                  string
  +--rw acl-type?                 acl-type

```

Figure 1





## **4. ACL YANG Models**

### **4.1. IETF Access Control List module**

"ietf-access-control-list" is the standard top level module for Access lists. The "access-lists" container stores a list of "acl". Each "acl" has information identifying the access list by a name("acl-name") and a list("access-list-entries") of rules associated with the "acl-name". Each of the entries in the list("access-list-entries"), indexed by the string "rule-name", has containers defining "matches" and "actions". The "matches" define criteria used to identify patterns in "ietf-packet-fields". The "actions" define behavior to undertake once a "match" has been identified.

```
<CODE BEGINS>file "ietf-access-control-list@2015-05-03.yang"
module ietf-access-control-list {
  yang-version 1;
  namespace "urn:ietf:params:xml:ns:yang:ietf-access-control-list";
  prefix acl;
  import ietf-yang-types {
    prefix yang;
  }
  import ietf-packet-fields {
    prefix packet-fields;
  }
  organization "IETF NETMOD (NETCONF Data Modeling Language)
    Working Group";
  contact
    "WG Web: http://tools.ietf.org/wg/netmod/
    WG List: netmod@ietf.org
    WG Chair: Juergen Schoenwaelder
    j.schoenwaelder@jacobs-university.de
    WG Chair: Tom Nadeau
    tnadeau@lucidvision.com
    Editor: Dean Bogdanovic
    deanb@juniper.net
    Editor: Kiran Agrahara Sreenivasa
    kkoushik@brocade.com
    Editor: Lisa Huang
    lyihuang@juniper.net
    Editor: Dana Blair
    dblair@cisco.com";
  description
    "This YANG module defines a component that describing the
    configuration of Access Control Lists (ACLs).
    Copyright (c) 2015 IETF Trust and the persons identified as
    the document authors. 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](http://trustee.ietf.org/license-info) of the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>).

This version of this YANG module is part of RFC XXXX; see the RFC itself for full legal notices.";

```
revision 2015-03-17 {
  description
    "Base model for Network Access Control List (ACL).";
  reference
    "RFC XXXX: Network Access Control List (ACL)
    YANG Data Model";
}
identity acl-base {
  description
    "Base Access Control List type for all Access Control List type
    identifiers.";
}
identity ip-acl {
  base acl:acl-base;
  description
    "IP Access Control List is a common name for lists that contain
    layer 3 and/or layer 4 match conditions.";
}
identity eth-acl {
  base acl:acl-base;
  description
    "Ethernet Access Control List is name for layer 2 Ethernet
    technology Access Control List types, like 10/100/1000baseT or
    WiFi Access Control List";
}
typedef acl-type {
  type identityref {
    base acl-base;
  }
  description
    "This type is used to refer to an Access Control List
    (ACL) type";
}
typedef access-control-list-ref {
  type leafref {
    path "/access-lists/acl/acl-name";
  }
  description
    "This type is used by data models that need to reference an
    Access Control List";
```



```
}
container access-lists {
  description
    "This is a top level container for Access Control Lists.
    It can have one or more Access Control Lists.";
  list acl {
    key "acl-name";
    description
      "An Access Control List(ACL) is an ordered list of
      Access List Entries (ACE). Each Access Control Entry has a
      list of match criteria and a list of actions.
      Since there are several kinds of Access Control Lists
      implemented with different attributes for
      different vendors, this
      model accommodates customizing Access Control Lists for
      each kind and for each vendor.";
    container acl-oper-data {
      config false;
      description
        "Overall Access Control List operational data";
    }
    container access-list-entries {
      description
        "The access-list-entries container contains
        a list of access-list-entries(ACE).";
      list ace {
        key "rule-name";
        ordered-by user;
        description
          "List of access list entries(ACE)";
        container matches {
          description
            "Definitions for match criteria for this Access List
            Entry.";
          choice ace-type {
            description
              "Type of access list entry.";
            case ace-ip {
              description "IP Access List Entry.";
              choice ace-ip-version {
                description
                  "IP version used in this Access List Entry.";
                case ace-ipv4 {
                  uses packet-fields:acl-ipv4-header-fields;
                }
                case ace-ipv6 {
                  uses packet-fields:acl-ipv6-header-fields;
                }
              }
            }
          }
        }
      }
    }
  }
}
```



```
    }
    uses packet-fields:acl-ip-header-fields;
  }
  case ace-eth {
    description
      "Ethernet Access List entry.";
    uses packet-fields:acl-eth-header-fields;
  }
}
uses packet-fields:metadata;
}
container actions {
  description
    "Definitions of action criteria for this Access List
Entry.";
  choice packet-handling {
    default "deny";
    description
      "Packet handling action.";
    case deny {
      leaf deny {
        type empty;
        description
          "Deny action.";
      }
    }
    case permit {
      leaf permit {
        type empty;
        description
          "Permit action.";
      }
    }
  }
}
}
container ace-oper-data {
  config false;
  description
    "Operational data for this Access List Entry.";
  leaf match-counter {
    type yang:counter64;
    description
      "Number of matches for this Access List Entry";
  }
}
}
leaf rule-name {
  type string;
  description
```





```

        "A unique name identifying this Access List
        Entry(ACE).";
    }
}
}
leaf acl-name {
    type string;
    description
        "The name of access-list. A device MAY restrict the length
and value of this name, possibly space and special
characters are not allowed.";
}
leaf acl-type {
    type acl-type;
    description
        "It is recommended to have an Access Control List with
uniform access list entries, all of the same type. When
this type is not explicitly specified, if vendor
implementation permits, the access control entries
in the list can be mixed,
by containing L2, L3 and L4 entries";
}
}
}
}
<CODE ENDS>

```

#### [4.2.](#) IETF-PACKET-FIELDS module

The packet fields module defines the necessary groups for matching on fields in the packet including ethernet, ipv4, ipv6, transport layer fields and metadata. Since the number of match criteria is very large, the base draft does not include these directly but references them by "uses" to keep the base module simple. In case more match conditions are needed, those can be added by augmenting choices within container "matches" in ietf-access-control-list.yang model

```

<CODE BEGINS>file "ietf-packet-fields@2015-06-11.yang"
module ietf-packet-fields {
    yang-version 1;
    namespace "urn:ietf:params:xml:ns:yang:ietf-packet-fields";
    prefix packet-fields;
    import ietf-inet-types {
        prefix inet;
    }
    import ietf-yang-types {
        prefix yang;
    }
}

```



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

contact

"WG Web: <http://tools.ietf.org/wg/netmod/>

WG List: [netmod@ietf.org](mailto:netmod@ietf.org)

WG Chair: Juergen Schoenwaelder

[j.schoenwaelder@jacobs-university.de](mailto:j.schoenwaelder@jacobs-university.de)

WG Chair: Tom Nadeau

[tnadeau@lucidvision.com](mailto:tnadeau@lucidvision.com)

Editor: Dean Bogdanovic

[deanb@juniper.net](mailto:deanb@juniper.net)

Editor: Kiran Agrahara Sreenivasa

[kkoushik@brocade.com](mailto:kkoushik@brocade.com)

Editor: Lisa Huang

[lyihuang@juniper.net](mailto:lyihuang@juniper.net)

Editor: Dana Blair

[dblair@cisco.com](mailto:dblair@cisco.com)";

description

"This YANG module defines groupings that are used by ietf-access-control-list YANG module. Their usage is not limited to ietf-access-control-list and can be used anywhere as applicable.

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. 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

(<http://trustee.ietf.org/license-info>).

This version of this YANG module is part of RFC XXXX; see the RFC itself for full legal notices.";

revision 2015-06-11 {

description

"Initial version of packet fields used by ietf-access-control-list";

reference

"RFC XXXX: Network Access Control List (ACL) YANG Data Model";

}

grouping acl-transport-header-fields {

description

"Transport header fields";

container source-port-range {

description

"Inclusive range representing source ports to be used.

When only lower-port is present, it represents a single port.";

leaf lower-port {



```
        type inet:port-number;
        mandatory true;
        description
            "Lower boundary for port.";
    }
    leaf upper-port {
        must ". >= ../lower-port" {
            error-message
                "The upper-port must be greater than or equal to lower-port";
        }
        type inet:port-number;
        description
            "Upper boundary for port . If existing, the upper port
            must be greater or equal to lower-port.";
    }
}
container destination-port-range {
    description
        "Inclusive range representing destination ports to be used. When
only lower-port is present, it represents a single port.";
    leaf lower-port {
        type inet:port-number;
        mandatory true;
        description
            "Lower boundary for port.";
    }
    leaf upper-port {
        must ". >= ../lower-port" {
            error-message
                "The upper-port must be greater than or equal to lower-port";
        }
        type inet:port-number;
        description
            "Upper boundary for port. If existing, the upper port must
            be greater or equal to lower-port";
    }
}
}
grouping acl-ip-header-fields {
    description
        "IP header fields common to ipv4 and ipv6";
    leaf dscp {
        type inet:dscp;
        description
            "Value of dscp.";
    }
    leaf protocol {
        type uint8;
```



```
    description
      "Internet Protocol number.";
  }
  uses acl-transport-header-fields;
}
grouping acl-ipv4-header-fields {
  description
    "Fields in IPv4 header.";
  leaf destination-ipv4-network {
    type inet:ipv4-prefix;
    description
      "Destination IPv4 address prefix.";
  }
  leaf source-ipv4-network {
    type inet:ipv4-prefix;
    description
      "Source IPv4 address prefix.";
  }
}
grouping acl-ipv6-header-fields {
  description
    "Fields in IPv6 header";
  leaf destination-ipv6-network {
    type inet:ipv6-prefix;
    description
      "Destination IPv6 address prefix.";
  }
  leaf source-ipv6-network {
    type inet:ipv6-prefix;
    description
      "Source IPv6 address prefix.";
  }
  leaf flow-label {
    type inet:ipv6-flow-label;
    description
      "IPv6 Flow label.";
  }
}
reference
  "RFC 4291: IP Version 6 Addressing Architecture
RFC 4007: IPv6 Scoped Address Architecture
RFC 5952: A Recommendation for IPv6 Address Text Representation";
}
grouping acl-eth-header-fields {
  description
    "Fields in Ethernet header.";
  leaf destination-mac-address {
    type yang:mac-address;
    description
```





```
        "Destination IEEE 802 MAC address.";
    }
    leaf destination-mac-address-mask {
        type yang:mac-address;
        description
            "Destination IEEE 802 MAC address mask.";
    }
    leaf source-mac-address {
        type yang:mac-address;
        description
            "Source IEEE 802 MAC address.";
    }
    leaf source-mac-address-mask {
        type yang:mac-address;
        description
            "Source IEEE 802 MAC address mask.";
    }
    reference
        "IEEE 802: IEEE Standard for Local and Metropolitan Area
        Networks: Overview and Architecture.";
}
grouping timerange {
    description
        "Time range contains time
        segments to allow access-control-list to be
        active/inactive when the system time
        is between the range.";
    container absolute-time {
        description
            "Absolute time and date that
            the associated function starts
            going into effect.";
        leaf start {
            type yang:date-and-time;
            description
                "Absolute start time and date";
        }
        leaf end {
            type yang:date-and-time;
            description
                "Absolute end time and date";
        }
    }
    leaf active {
        type boolean;
        default "true";
        description
            "This object indicates whether the
            the ACL will be active(true) or
```



```
        inactive(false) during this time range.";
    }
}
grouping metadata {
    description
        "Fields associated with a packet which are not in
        the header.";
    leaf input-interface {
        type string;
        description
            "Packet was received on this interface.";
    }
    uses timerange;
}
}
<CODE ENDS>
```

#### [4.3.](#) An ACL Example

Requirement: Deny All traffic from 10.10.10.1 bound for host 10.10.10.255 from leaving.

In order to achieve the requirement, an name Access Control List is needed. The acl and aces can be described in CLI as the following:

```
access-list ip sample-ip-acl
    deny tcp host 10.10.10.1 host 10.10.10.255
```

Here is the example acl configuration xml:



```
<rpc message-id="101" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0">
<edit-config>
  <target>
    <running/>
  </target>
  <config>
    <access-lists "urn:ietf:params:xml:ns:yang:ietf-acl:1.0">
      <acl>
        <acl-name>sample-ip-acl</acl-name>
        <access-list-entries>
          <ace>
            <rule-name>rule1</rule-name>
            <matches>
              <destination-ipv4-network>
                10.10.10.255/24
              </destination-ipv4-network>
              <source-ipv4-network>
                10.10.10.1/24
              </source-ipv4-network>
            </matches>
            <actions>
              <deny/>
            </actions>
          </ace>
        </access-list-entries>
      </acl>
    </access-lists>
  </config>
</edit-config>
</rpc>
```

#### [4.4.](#) Port Range Usage Example

When a lower-port and an upper-port are both present, it represents a range between lower-port and upper-port with both the lower-port and upper-port are included. When only a lower-port presents, it represents a single port.

With the follow XML snippet:

```
<source-port-range>
  <lower-port>16384</lower-port>
  <upper-port>16387</upper-port>
</source-port-range>
```

This represents source ports 16384,16385, 16386, and 16387.



With the follow XML snippet:

```
<source-port-range>
  <lower-port>16384</lower-port>
  <upper-port>65535</upper-port>
</source-port-range>
```

This represents source ports greater than/equal to 16384.

With the follow XML snippet:

```
<source-port-range>
  <lower-port>21</lower-port>
</source-port-range>
```

This represents port 21.

## 5. Linux nftables

As Linux platform is becoming more popular as networking platform, the Linux data model is changing. Previously ACLs in Linux were highly protocol specific and different utilities were used (iptables, ip6tables, arptables, ebtables), so each one had separate data model. Recently, this has changed and a single utility, nftables, has been developed. With a single application, it has a single data model for firewall filters and it follows very similarly to the ietf-access-control list module proposed in this draft. The nftables support input and output ACEs and each ACE can be defined with match and action.

In the example below, it shows nftable configuration that accepts and count packets. It contains a

```
table ip filter {
  chain output {
    type filter hook output priority 0;
    counter packets 1 bytes 84 accept
  }
}
```

There are many similarities between Linux nftables and IETF ACL YANG data models. It should be fairly easy to do translation between ACL YANG model described in this draft and Linux nftables.





## 6. Security Considerations

The YANG module defined in this memo is designed to be accessed via the NETCONF protocol [[RFC6241](#)] [[RFC6241](#)]. The lowest NETCONF layer is the secure transport layer and the mandatory-to-implement secure transport is SSH [[RFC6242](#)] [[RFC6242](#)]. The NETCONF access control model [[RFC6536](#)] [[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 the 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:

/access-lists/acl/access-list-entries: This list specifies all the configured access list entries on the device. Unauthorized write access to this list can allow intruders to access and control the system. Unauthorized read access to this list can allow intruders to spoof packets with authorized addresses thereby compromising the system.

## 7. IANA Considerations

This document registers a URI in the IETF XML registry [[RFC3688](#)] [[RFC3688](#)]. Following the format in [RFC 3688](#), the following registration is requested to be made:

URI: urn:ietf:params:xml:ns:yang:ietf-access-control-list

URI: urn:ietf:params:xml:ns:yang:ietf-packet-fields

Registrant Contact: The IESG.

XML: N/A, the requested URI is an XML namespace.

This document registers a YANG module in the YANG Module Names registry [[RFC6020](#)].

name: ietf-access-control-list namespace:

urn:ietf:params:xml:ns:yang:ietf-access-control-list prefix: ietf-acl

reference: RFC XXXX



name: ietf-packet-fields namespace: urn:ietf:params:xml:ns:yang:ietf-packet-fields prefix: ietf-packet-fields reference: RFC XXXX

## **8. Acknowledgements**

Alex Clemm, Andy Bierman and Lisa Huang started it by sketching out an initial IETF draft in several past IETF meetings. That draft included an ACL YANG model structure and a rich set of match filters, and acknowledged contributions by Louis Fourie, Dana Blair, Tula Kraiser, Patrick Gili, George Serpa, Martin Bjorklund, Kent Watsen, and Phil Shafer. Many people have reviewed the various earlier drafts that made the draft went into IETF charter.

Dean Bogdanovic, Kiran Agrahara Sreenivasa, Lisa Huang, and Dana Blair each evaluated the YANG model in previous draft separately and then work together, to created a new ACL draft that can be supported by different vendors. The new draft removes vendor specific features, and gives examples to allow vendors to extend in their own proprietary ACL. The earlier draft was superseded with the new one that received more participation from many vendors.

## **9. References**

### **9.1. Normative References**

- [RFC3688] Mealling, M., "The IETF XML Registry", [BCP 81](#), [RFC 3688](#), January 2004.
- [RFC6020] Bjorklund, M., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", [RFC 6020](#), October 2010.
- [RFC6241] Enns, R., Bjorklund, M., Schoenwaelder, J., and A. Bierman, "Network Configuration Protocol (NETCONF)", [RFC 6241](#), June 2011.
- [RFC6242] Wasserman, M., "Using the NETCONF Protocol over Secure Shell (SSH)", [RFC 6242](#), June 2011.
- [RFC6536] Bierman, A. and M. Bjorklund, "Network Configuration Protocol (NETCONF) Access Control Model", [RFC 6536](#), March 2012.

### **9.2. Informative References**

- [RFC5101] Claise, B., "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information", [RFC 5101](#), January 2008.



## [Appendix A](#). Extending ACL model examples

### [A.1](#). Example of extending existing model for route filtering

With proposed modular design, it is easy to extend the model with other features. Those features can be standard features, like route filters. Route filters match on specific IP addresses or ranges of prefixes. Much like ACLs, they include some match criteria and corresponding match action(s). For that reason, it is very simple to extend existing ACL model with route filtering. The combination of a route prefix and prefix length along with the type of match determines how route filters are evaluated against incoming routes. Different vendors have different match types and in this model we are using only ones that are common across all vendors participating in this draft. As in this example, the base ACL model can be extended with company proprietary extensions, described in the next section.

```
file "example-ext-route-filter@2015-02-14.yang"
```

```
module example-ext-route-filter {
  yang-version 1;
  namespace "urn:ietf:params:xml:ns:yang:example-ext-route-filter";
  prefix example-ext-route-filter;
  import ietf-inet-types {
    prefix "inet";
  }
  import ietf-access-control-list {
    prefix "ietf-acl";
  }
  organization
    "Route modele group.";

  contact
    "abc@abc.com";

  description "
    This module describes route filter as a collection of
    match prefixes. When specifying a match prefix, you
    can specify an exact match with a particular route or
    a less precise match. You can configure either a
    common action that applies to the entire list or an
    action associated with each prefix.
  ";

  revision 2015-05-03 {
    description
      "Creating Route-Filter extension model based on
      ietf-access-control-list model";
    reference " ";
  }
}
```



```
}
augment "/ietf-acl:access-lists/ietf-acl:acl/
  ietf-acl:access-list-entries/ietf-acl:ace/ietf-acl:matches"{
  description "
    This module augments the matches container in the ietf-acl
    module with route filter specific actions
  ";
  choice route-prefix{
    description "Define route filter match criteria";
    case range {
      description
        " Route falls between the lower prefix/prefix-length
        and the upperprefix/prefix-length.";
      choice ipv4-range {
        description "Defines the IPv4 prefix range";
        leaf v4-lower-bound {
          type inet:ipv4-prefix;
          description
            "Defines the lower IPv4 prefix/prefix length";
        }
        leaf v4-upper-bound {
          type inet:ipv4-prefix;
          description
            "Defines the upper IPv4 prefix/prefix length";
        }
      }
      choice ipv6-range {
        description "Defines the IPv6 prefix/prefix range";
        leaf v6-lower-bound {
          type inet:ipv6-prefix;
          description
            "Defines the lower IPv6 prefix/prefix length";
        }
        leaf v6-upper-bound {
          type inet:ipv6-prefix;
          description
            "Defines the upper IPv6 prefix/prefix length";
        }
      }
    }
  }
}
```





## [A.2.](#) A company proprietary module example

Module "example-newco-acl" is an example of company proprietary model that augments "ietf-acl" module. It shows how to use 'augment' with an XPath expression to add additional match criteria, action criteria, and default actions when no ACE matches found. All these are company proprietary extensions or system feature extensions. "example-newco-acl" is just an example and it is expected from vendors to create their own proprietary models.

The following figure is the tree structure of example-newco-acl. In this example, /ietf-acl:access-lists/ietf-acl:acl/ietf-acl:access-list-entries/ ietf-acl:ace/ietf-acl:matches are augmented with a new choice, protocol-payload-choice. The protocol-payload-choice uses a grouping with an enumeration of all supported protocol values. In other example, /ietf-acl:access-lists/ietf-acl:acl/ietf-acl:access-list-entries/ ietf-acl:ace/ietf-acl:actions are augmented with new choice of actions.

```
module: example-newco-acl
  augment /ietf-acl:access-lists/ietf-acl:acl/
    ietf-acl:access-list-entries/ietf-acl:ace/ietf-acl:matches:
      +--rw (protocol-payload-choice)?
        +--:(protocol-payload)
          +--rw protocol-payload* [value-keyword]
            +--rw value-keyword    enumeration
  augment /ietf-acl:access-lists/ietf-acl:acl/
    ietf-acl:access-list-entries/ietf-acl:ace/ietf-acl:actions:
      +--rw (action)?
        +--:(count)
          | +--rw count?          string
        +--:(policer)
          | +--rw policer?        string
        +--:(hierarchical-policer)
          +--rw hierarchitacl-policer? string
  augment /ietf-acl:access-lists/ietf-acl:acl:
    +--rw default-actions
    +--rw deny?    empty
```

file "newco-acl@2015-03-04.yang"

```
module example-newco-acl {
  yang-version 1;

  namespace "urn:newco:params:xml:ns:yang:example-newco-acl";

  prefix example-newco-acl;
```



```
import ietf-acl {
  prefix "ietf-acl";
}

revision 2015-05-03{
  description "Creating NewCo proprietary extensions to ietf-acl model";
}

augment "/ietf-acl:access-lists/ietf-acl:access-list
/ietf-acl:access-list-entries/
  ietf-acl:access-list-entry/ietf-acl:matches" {
  description "Newco proprietary simple filter matches";
  choice protocol-payload-choice {
    list protocol-payload {
      key value-keyword;
      ordered-by user;
      description "Match protocol payload";
      uses match-simple-payload-protocol-value;
    }
  }
}

augment "/ietf-acl:access-lists/ietf-acl:access-list/
  ietf-acl:access-list-entries/ietf-acl:access-list-entry/
  ietf-acl:actions" {
  description "Newco proprietary simple filter actions";
  choice action {
    case count {
      description "Count the packet in the named counter";
      leaf count {
        type string;
      }
    }
    case policer {
      description "Name of policer to use to rate-limit traffic";
      leaf policer {
        type string;
      }
    }
    case hierarchical-policer {
      description "Name of hierarchical policer to use to
rate-limit traffic";
      leaf hierarchitac1-policer{
        type string;
      }
    }
  }
}
```



```
augment "/ietf-acl:access-lists/ietf-acl:access-list" {
  container default-actions {
    description "Actions that occur if no access-list entry is matched.";
    leaf deny {
      type empty;
    }
  }
}

grouping match-simple-payload-protocol-value {
  leaf value-keyword {
    description "(null)";
    type enumeration {
      enum icmp {
        description "Internet Control Message Protocol";
      }
      enum icmp6 {
        description "Internet Control Message Protocol Version 6";
      }
      enum range {
        description "Range of values";
      }
    }
  }
}
```

Draft authors expect that different vendors will provide their own yang models as in the example above, which is the augmentation of the base model

### **A.3. Attaching Access Control List to interfaces**

Access control list typically does not exist in isolation. Instead, they are associated with a certain scope in which they are applied, for example, an interface of a set of interfaces. How to attach an access control list to an interface (or other system artifact) is outside the scope of this model, as it depends on the specifics of the system model that is being applied. However, in general, the general design pattern will involved adding a data node with a reference, or set of references, to ACLs that are to be applied to the interface. For this purpose, the type definition "access-control-list-ref" can be used.

This is an example of attaching an Access Control List to an interface.



```
import ietf-access-control-list {
  prefix "ietf-acl";
}
import ietf-interface {
  prefix "ietf-if";
}
import ietf-yang-types {
  prefix "yang";
}

augment "/ietf-if:interfaces/ietf-if:interface" {
  description "Apply ACL to interfaces";
  container acl{
    description "ACL related properties.";
    leaf acl-name {
      type ietf-acl:acl-ref;
      mandatory true;
      description "Access Control List name.";
    }
    leaf match-counter {
      type yang:counter64;
      config false;
      description
        "Total match count for Access Control
        List on this interface";
    }
    choice direction {
      leaf in { type empty;}
      leaf out { type empty;}
    }
  }
}

augment "/ietf-acl:access-lists/ietf-acl:acl/ietf-acl:acl-oper-data" {
  description
    "This is an example on how to apply acl to a target to collect
    operational data";
  container targets{
    choice interface{
      leaf-list interface-name{
        type ietf-if:interface-ref;
      }
    }
  }
}
```





Authors' Addresses

Dean Bogdanovic  
Juniper Networks

Email: deanb@juniper.net

Kiran Agrahara Sreenivasa  
Brocade Communications System

Email: kkoushik@brocade.com

Lisa Huang  
Juniper Networks

Email: lyihuang@juniper.net

Dana Blair  
Cisco Systems

Email: dblair@cisco.com

