Revision Letter

Editor: Jaehoon Paul Jeong
Date: February 11, 2022

OLD: draft-ietf-i2nsf-nsf-facing-interface-dm-20
NEW: draft-ietf-i2nsf-nsf-facing-interface-dm-21

Dear Tom Petch,

I sincerely appreciate your detailed comments to improve our NSF-Facing Interface YANG Data Model. I have addressed all your comments one by one. I use bold font for your comments and use a regular blue font for my responses with the prefix "=> [PAUL]".

With an IESG telechat scheduled for 17th February, I believe that -17 goes in the wrong direction, with its expanded descriptions, with its import of packet filters and its addition of application-protocol identity.

The I-D has been reviewed many times without comment on the descriptions so when one reviewer does comment I see that as more a comment on the reviewer than on the I-D:-) Here it is problematic because the same YANG definitions appear in other I-D without the expanded descriptions so there is a mismatch within the set of I-D. And as an AD said recently, you should reference and not replicate technical matter (which as ever highlights the lack of a common document for all the things that are common although I do not advocate tackling that at this stage of the cycle; perhaps a reissue of the framework document, RFC8329, or an update of the terminology one) so if expanded descriptions are needed they should be in one place. The expanded descriptions will need some editing of the English (perhaps a lot, IMHO) but since I do not think that they should exist I will not comment further on that.
=> [PAUL] One side of the argument (by Joe Clarke) says that it is better to have descriptive text to help a reader understand without having to jump between documents. We authors will update the identities in the documents to have the same descriptions for the shared identities in the I2NSF YANG data model documents.

Packet filters (RFC8519) were referenced in the framework document and if they were in this I-D from the start then I might have been ok with that but I see such a big change so late in the day as the wrong thing to do especially as I do not see packet filters as the right approach. They are clumsy. Many YANG modules want to specify a range or a single value, of IP address and such like and many if not most achieve that with 'min' and 'max' with 'min = max' for the single value (some use the absence of max to denote this but for me that is fail danger). With that approach the model has two leaves
min
max

With the packet filter approach you have e.g.

```
   grouping port-range-or-operator {
     choice port-range-or-operator {
       case range {
         leaf lower-port {
         leaf upper-port {
       case operator {
         <eq neq gte lte>
         leaf operator {
         leaf port {
```

which I find overly complex and so prone to misunderstanding and error.

=> It is a good point about the usage of IP addresses that are sometimes necessary to specify a range of IP addresses. We include the usage of the "min" and "max" model ("start" and "end" in the data model) by augmenting the “acl-ipv4-header-fields” node with a case of ipv4-range for IPv4 and IPv6 Addresses. This model should allow a user to configure IP addresses in two ways, i.e., with IP prefix and IP range. The update in the data model is as follows:

NEW:

```
        container ipv4 {
          [...]
          uses packet-fields:acl-ipv4-header-fields {
            augment destination-network {
              case destination-ipv4-range {
                list destination-ipv4-range {
                  key "start end";
                  uses ipv4-range;
                  description
                    "The list of IPv4 addresses specified with a
                     start IPv4 address and an end IPv4 address.
                     If only one value is needed, then set both
                     start and end to the same value.
                     Note that the 'end' IPv4 address MUST be equal
                     to or greater than the 'start' IPv4 address.";
                }
              }
            }
            augment source-network {
              case source-ipv4-range {
                list source-ipv4-range {
                  key "start end";
                  uses ipv4-range;
                  description
                    "The list of IPv4 addresses specified with a
                     start IPv4 address and an end IPv4 address.
                     If only one value is needed, then set both
                     start and end to the same value.
                     Note that the 'end' IPv4 address MUST be equal
                     to or greater than the 'start' IPv4 address.";
                }
              }
            }
          }
        }
```

```
        container ipv6 {
          [...]

          uses packet-fields:acl-ip-header-fields;
          uses packet-fields:acl-ipv6-header-fields {
            augment destination-network {
              case destination-ipv6-range {
                list destination-ipv6-range {
                  key "start end";
                  uses ipv6-range;
                  description
                    "The list of IPv6 addresses specified with a
                     start IPv6 address and an end IPv6 address.
                     If only one value is needed, then set both
                     start and end to the same value.
                     Note that the 'end' IPv6 address MUST be equal
                     or greater than the 'start' IPv6 address.";
                }
              }
            }
            augment source-network {
              case source-ipv6-range {
                list source-ipv6-range {
                  key "start end";
                  uses ipv6-range;
                  description
                    "The list of IPv6 addresses specified with a
                     start IPv6 address and an end IPv6 address.
                     If only one value is needed, then set both
                     start and end to the same value.
                     Note that the 'end' IPv6 address MUST be equal
                     or greater than the 'start' IPv6 address.";
                }
              }
            }
          }
        }
```

For the port number of the transport protocol, the usage of “operator” may be useful to simplify the configuration. For example, the “neq” (not equal) can be used to define one single port number that is not needed as follows:

```
 <i2nsf-security-policy
xmlns="urn:ietf:params:xml:ns:yang:ietf-i2
nsf-policy-rule-for-nsf">
 <name>flood_attack_mitigation</name>
 <rules>

<name>mitigate_http_and_https_flood_attack
</name>
  <condition>
   <ipv4>
```

```
<i2nsf-security-policy
xmlns="urn:ietf:params:xml:ns:yang:ietf-i2
nsf-policy-rule-for-nsf">
 <name>flood_attack_mitigation</name>
 <rules>

<name>mitigate_http_and_https_flood_attack
</name>
  <condition>
   <ipv4>
```

```
<destination-ipv4-network>192.0.2.0/24</de
stination-ipv4-network>
   </ipv4>
   <tcp>
    <destination-port-number>
     <operator>neq</operator>
     <port>21</port>
    </destination-port-number>
   </tcp>
  </condition>
  <action>
   <advanced-action>
    <attack-mitigation-control>
      anti-ddos
    </attack-mitigation-control>
   </advanced-action>
  </action>
 </rules>
</i2nsf-security-policy>
```

```
<destination-ipv4-network>192.0.2.0/24</de
stination-ipv4-network>
   </ipv4>
   <tcp>
    <destination-port-number>
     <port-numbers>
      <start>0</start>
      <end>20</end>
     </port-numbers>
     <port-numbers>
      <start>22</start>
      <end>65535</end>
     </port-numbers>
    </destination-port-number>
   </tcp>
  </condition>
  <action>
   <advanced-action>
    <attack-mitigation-control>
      anti-ddos
    </attack-mitigation-control>
   </advanced-action>
  </action>
 </rules>
</i2nsf-security-policy>
```

'application-protocol' has made an appearance in -17 and I do not know where that came from. I can see the need for applications in 'consumer-facing' but it seemed of little relevance in 'nsf-facing' with its emphasis on ethernet and such like so the difference between the I-D seemed logical to me. And in incorporating the YANG identity, with references, you have, as ever, failed to add the references to the I-D References introducing another ten errors.

=> [PAUL] The “identity application-protocol” is used to configure a packet filter based on application protocols.

Some references are added for the NSF-Facing Interface Document. The added references are as follows:

- draft-ietf-httpbis-semantics (replacing RFC 2818, 7230, and 7231)
- draft-ietf-httpbis-messaging (replacing RFC 2818, 7230, and 7231)
- RFC 854
- RFC 959
- RFC 1939
- RFC 2595
- RFC 4250
- RFC 5321
- RFC 9051

I note the shortening of the names which can be a good idea if it were done consistently across the I-D and it were done at an earlier stage. (I note that the examples would appear to be in line with this; on an earlier occasion they were not). In places, though it may have gone too

far; sometimes there are too many fields with an identifier of 'name' IMHO and a prefix thereto would be helpful. And as ever this introduces inconsistencies across the set of I-D which should be found and fixed, not an exercise I am likely to undertake any time soon. And as and when the terminology diverges from RFC8329 then I think some comment thereon is called for.

=> [PAUL] As you mentioned, it seems reasonable to have a prefix where there is more than one 'name'. We updated the naming as follows:

OLD:

```
  list groups {
    key "name";
    description
      "This is a group for rules";

    leaf name {
      type string;
      description
        "This is the name of the group for rules";
    }

    leaf-list rule-name {
      type leafref {
        path
          "../../../rules/name";
      }
      description
        "The names of the rules to be grouped.";
    }
```

NEW:

```
  list groups {
    key "group-name";
    description
      "This is a group for rules";

    leaf group-name {
      type string;
      description
        "This is the name of the group for rules";
    }

    leaf-list rule-name {
      type leafref {
        path
          "../../../rules/name";
      }
      description
        "The names of the rules to be grouped.";
    }
```

I realise that multiple versions of nsf-facing have appeared since -17 but I planned my work to be complete in time for the IETF telechat and never imagined that there would be such extensive changes so late in the day, I have yet to look at them. I may, I may not.

Tom Petch

Thanks for your sincere help and support.

Best Regards,
Jaehoon (Paul) Jeong