RTG Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 9, 2020

A. Clemm
P. Pillay-Esnault
U. Chunduri
Futurewei

July 8, 2019

Preferred Path Routing (PPR) OAM and Accounting
draft-cpc-rtgwg-ppr-oam-00

Abstract

This document defines OAM and traffic accounting capabilities for Preferred Path Routing (PPR) for IS-IS and OSPF protocols. Specifically, this document specifies OAM capabilities that allow to assert proper PPR connectivity and to trace PPR path information. In addition, a set of statistics and operational data to facilitate PPR traffic accounting on a per-PPR path basis are defined. This includes a number of Information Elements that extend IPFIX to export path information, as well as a YANG Data Model to be used in conjunction with management and control protocols. Collectively the capabilities defined in this document provide network operators with the necessary means to ensure proper working of their PPR deployments.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC2119 [RFC2119], RFC8174 [RFC8174] when, and only when they appear in all capitals, as shown here.

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 January 9, 2020.

Copyright Notice

Table of Contents

1. Introduction 3
2. Key Words 3
3. Definition and Acronyms 4
4. PPR Ping and Trace Functionality 4
4.1. PPR Trace Functionality in Strict Mode 6
4.2. PPR Trace Functionality in Loose Mode 7
5. Traffic Accounting through IGP PPR-Attribute Sub-TLVs 8
6. Path Statistics and IP[F/P]IX 9
7. A YANG Data Model for PPR Monitoring 11
7.1. Motivation and Overview 11
7.2. YANG Data Model 13
8. Acknowledgements 16
9. IANA Considerations 16
9.1. IGP Path Attributes 16
9.2. IPFIX Information Elements 16
9.3. YANG Data Model 16
10. Security Considerations 17
10.1. Path Trace and Ping 17
10.2. IGPs and IPFIX 17
10.3. YANG Data Model 18
11. References 18
11.1. Normative References 18
11.2. Informative References 20
Authors' Addresses 21

1. Introduction

Preferred Path Routing (PPR) allows to route packets along a custom path on the basis of a path identifier (PPR-ID) as opposed to individual segments in the packet. Interior Gateway Protocols (IGPs) nodes compute nexthops based on the path description for the prefix and take proper forwarding actions for the paths that they are a part of. PPR may be used with different data planes, e.g. IPv4, IPv6, MPLS, SRv6. Dissemination and nexthop computation of path information is described in [I-D.chunduri-lsr-isis-preferred-path-routing] and [I-D.chunduri-lsr-ospf-preferred-path-routing].

In order to operate, administer, and maintain PPR deployments, network operators must be given tools that allow them to ensure that PPR is working as expected and that allow them to troubleshoot any potential problems. This includes assessing whether remote destinations can indeed be reached as intended using a given preferred path, and to verify path elements along the path being taken. Traditionally, ping (ICMP echo request on IP networks, LSP ping in MPLS networks) and traceroute operations are used for those purposes. In order to facilitate operation of PPR, analogous capabilities with PPR-specific extensions are useful which are defined in this document.

In addition, for purposes of traffic accounting and ongoing operations, it can be very useful to maintain certain PPR statistics. This allows, for example, to assess times and volume when traffic over preferred paths is occurring. This document defines new PPR path attributes as defined in Section 5 and this can be optionally signaled for the preferred paths selectively to account for the traffic on every path segment (where critical SLAs are needed). IPFIX is commonly used to maintain and export flow-specific information from any node in the network. This document introduces a number of new IPFIX Information Elements that are useful in conjunction with PPR.

Some of the same statistics maintained on a per-flow basis may be useful to maintain not just for the duration of a flow, but for the lifetime of the path. That information is considered part of regular management information and subject to specification as part of a YANG data model.

2. Key Words

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.

3. Definition and Acronyms

- IE: IPFIX Information Element
- IGP: Interior Gateway Protocol
- IPFIX: IP Flow Information eXport
- IS-IS Link State PDU
- PPR: Preferred Path Routing/Route
- PPR-ID: Preferred Path Route Identifier, a data plane identifier
- SRH: Segment Routing Header - IPv6 routing Extension header
- SRv6: Segment Routing with Ipv6 data plane with SRH
- TE: Traffic Engineering

4. PPR Ping and Trace Functionality

Ping and Trace capabilities depend on the underlying data plane being used for the PPR path. Different mechanisms exist for each of those data planes. As PPR support does not require any change in the existing dataplane, for each of those cases, existing ping and trace core mechanisms continue to work seamlessly without modification:

IPv4 data plane: ICMP [RFC792]

IPv6 data plane: ICMPv6 [RFC4443]

MPLS data plane: LSP Ping [RFC8287]

SRv6 data plane: Mechanism under development

To manage PPR, network operators may require additional information that is specific to PPR. For example, PPR may be installed by different IGPs(OSPF/ISIS) and supports two types of paths:

(1) Strict: where all the nodes on the path are specified and the traffic is forwarded at every single hop to the next node defined in the path.

(2) Loose: where not all nodes are described in the path. There are section(s) of the path that are unspecified and the packets are forwarded based on the local router forwarding until the next node in the path description.

The current mechanisms need to be enhanced to carry the PPR specific mechanism but otherwise it has minimal changes. In this section, a simple topology is shown in Figure 1 is used to illustrate the different traces of PPR paths described in the subsequent sub sections below.

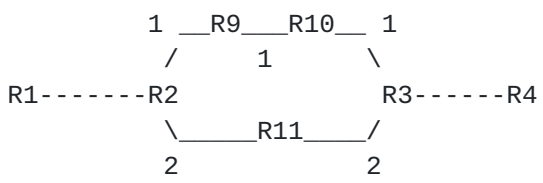

Figure 1. Topology example

PPR paths from topology
PPR-ID: (strict) -> R1, R2, R3, R11, R4

PPR-ID: (Loose) -> R1, R2, R3, R4
The section between R2-R3 is loose.
The best path is via R9 and 10. Path is R1, R2, (R9, R10), R3, R4

Node	Node IPaddr	SPF Cost to R3
R1	1.1.1.1	
R2	2.2.2.2	costs via R9+R10=3, via R11=4
R3	3.3.3.3	
R4	4.4.4.4	
R9	9.9.9.9	
R10	10.10.10.10	
R11	11.11.11.11	

The classic ping to a destination relying on the forwarding path works seamlessly as the PPR path is installed via IGPs. The PPR-ID is the same as the destination address/label in different dataplanes, therefore the existing ping mechanism remains unchanged in a network supporting PPR.

However, the traceroute output should display more PPR specific information such as whether or not the path from a node is loose or strict, or the source of the protocol that installed a PPR-ID or the dataplane information. In order to support PPR-specific trace information, a few additional capabilities are needed.

A PPR Trace is initiated from a node, using the PPR-ID as a parameter. The trace will then traverse the PPR path of the specified PPR-ID. In addition to the information that would be returned with a regular trace, PPR-specific information is returned from PPR nodes that are encountered along the path. For each PPR node, the following information is returned:

(1) Node Information: Provides the loopback identifier of the node (TBD).

(2) Loose-or-strict-or-none: Indicates whether the path from that node is loose or strict, or whether the node is PPR-capable but not one of the nodes designated in the path. The latter can be the case for transit nodes on a loose segment of the path.

(3) PPR-Origin Protocol: Indicates the protocol that installed the PPR-ID in the FIB.

(4) PPR-ID: Indicates the PPR-ID of the packet as received. This is needed in case of a loose path, leading to encapsulated PPR-ID, described further in Section 4.2.

These PPR specific information will be carried in ICMP data. These extensions are TBD for the various dataplanes.

4.1. PPR Trace Functionality in Strict Mode

The PPR strict mode indicates the PPR path on a hop by hop basis. As the PPR-ID represents a strict path to be followed, the ingress node should increment the TTL for the PPR-ID until all the egress node has returned all individual intermediate segments completely traced.

The PPR specific enhancements additional information are returned as part of trace information. The strict path to be trace is PPR-ID=4.4.4.5 represents the path (R1, R2, R11, R3, R4). An example of the output for traceroute for PPR-IS in IPV4 dataplane is given below:

```
>traceroute 4.4.4.5

traceroute to 4.4.4.5 (192.4.4.5), 30 hops max, 40 byte packets
1  1.1.1.1 (192.1.1.1)  6.819 ms  1.370 ms  1.281 ms
   PPR: ID:4.4.4.5 Mode=Strict Origin=IS-IS

2  2.2.2.2 (192.1.1.2)  4.437 ms  1.987 ms  2.335 ms
   PPR: ID:4.4.4.5 Mode=Strict Origin=IS-IS

3  11.11.11.11 (192.11.11.11)  9.830 ms  11.696 ms  5.478 ms
   PPR: ID:4.4.4.5 Mode=Strict Origin=IS-IS

4  3.3.3.3 (192.3.3.3)  9.448 ms  5.096 ms  7.518 ms
   PPR: ID:4.4.4.5 Mode=Strict Origin=IS-IS

5  4.4.4.4 (192.3.3.4)  9.448 ms  5.096 ms  7.518 ms
   PPR: ID:4.4.4.5 Mode=Strict Origin=IS-IS
```

Figure 2 Enhanced traceroute Output example

The output for IPv6 and MPLS traceroute should be augmented to display the PPR-ID specific parameters as shown above.

4.2. PPR Trace Functionality in Loose Mode

The PPR loose path functionality allows greater flexibility by letting the IGP calculate the best path on a loose section of the path described for PPR-ID. The traditional traceroute requires additional enhancement for the PPR trace to perform correctly . Per Figure 1, the loose path example is R1, R2,(R9, R10), R3, R4 where () represents the loose section of the path. An example of the output for traceroute for PPR-IS in IPV4 dataplane is given below:

```
>traceroute 4.4.4.5

traceroute to 4.4.4.5 (192.4.4.5), 30 hops max, 40 byte packets

1  1.1.1.1 (192.1.1.1)  6.819 ms  1.370 ms  1.281 ms
   PPR: ID:4.4.4.5 Mode=Strict Origin=OSPF

2  2.2.2.2 (192.1.1.2)  4.437 ms  1.987 ms  2.335 ms
   PPR: ID:4.4.4.5 Mode=Loose Origin=OSPF Multipaths=0

       1 9.9.9.9 (192.9.9.9)  9.830 ms  11.696 ms  5.478 ms
         PPR: ID:4.4.4.5 Mode=None Origin=OSPF

       2 10.10.10.10 (192.9.9.10)  4.230 ms  4.633 ms  5.789 ms
         PPR: ID:4.4.4.5 Mode=None Origin=OSPF

3  3.3.3.3 (192.3.3.3)  9.448 ms  5.096 ms  7.518 ms
   PPR: ID:4.4.4.5 Mode=Strict Origin =OSPF

4  4.4.4.4 (192.3.3.4)  9.448 ms  5.096 ms  7.518 ms
   PPR: ID:4.4.4.5 Mode=Strict Origin=OSPF
```

Figure 3 Enhanced traceroute Output example

The output for IPv6 and MPLS traceroute should be augmented to display the PPR-ID, Mode and Origin.

As the loose path functionality does not preclude the presence of equal cost multiple paths (ECMPs), the well documented limitations and solutions for classic traceroute applies here as well. For operational purposes, it might be of value to list all the multipaths. To acheive this the node just before the loose section MAY initiate a recursive traceroute to aggregate that information and send it with their ICMP Echo Reply message.

5. Traffic Accounting through IGP PPR-Attribute Sub-TLVs

Traffic for certain PPRs may have more stringent requirement w.r.t accounting for critical SLAs (e.g. 5G non-eMBB slice) and should account for any link/node failures along the path. Presence of "Packet Traffic Accounting" and "Traffic Statistics" Sub-TLVs below in IGP PPR-TLV instructs all the respective nodes along the path to provision the hardware and to account for the respective traffic statistics. Traffic accounting should happen, when the actual data traffic hits for the PPR-ID in the forwarding plane. This capability allows more granular and dynamic enablement of traffic statistics for only certain PPRs as needed.

As instruction for creating and deleting the traffic accounting for PPRs happen through IGP message processing, respective IGP's control plane security (Section 10) options are applicable to the PPR-TLV and Sub-TLVs thereof.

PPR-Attribute Sub-TLVs describe the attributes of that particular path. The following path attribute Sub-TLVs are defined for respective IGP PPR-TLV [I-D.chunduri-lsr-isis-preferred-path-routing] and [I-D.chunduri-lsr-ospf-preferred-path-routing].

- Type 10 (Suggested Value - IANA TBD): Packet Traffic accounting Sub-TLV. Length 0 and has no value field. Specifies to create a counter to count number of packets forwarded on this PPR-ID on each node in the path description.

- Type 11 (Suggested Value - IANA TBD): Traffic statistics in Bytes Sub-TLV. Length 0 and has no value field. Specifies to create a counter to count number of bytes forwarded on this PPR-ID specified in the network header (e.g. IPv4, IPv6) on each node in the path description.

How the accumulated traffic accounting information is distributed to a central entity is out of scope of this document. One can use any method (e.g. RESTCONF, Yang PUSH, gRPC) to extract the PPR-ID traffic accounting information from various nodes along the path.

6. Path Statistics and IP[F/P]IX

Network providers will appreciate the ability to collect certain statistics about PPR path usage, including how much traffic a PPR path carries and at what times from any node in the network. Such statistics can be useful to account for the degree of usage of a path and provide additional operational insights, including (for example) usage patterns and trending information.

One mechanism that is traditionally used to collect traffic statistics is IPFIX (IP Flow Information eXport [RFC7011]. IPFIX supports a very rich set of Information Elements containing various statistics, far more in fact than will be required for PPR path statistics. However, those statistics are collected only a per-flow basis.

In order to be able to collect statistics on a per-path basis, a set of new IPFIX Information Elements need to be defined that allow to capture a PPR-ID. In addition, these new IPFIX Information Elements need to be able to serve as a flow key. This allows separate IPFIX cache entries to be created on a per-path basis, and to export the corresponding records as IPFIX records. Of course, the records

exported do not refer to flows but to paths - IPFIX thus becomes de-facto extended to become IPPIX - IP Path Information eXport.

Flow records that have a PPR-ID as flow key SHALL be terminated and exported in the following events and/or per configurable policy:

- When no packet has not been observed on a path for a time interval defined by a flow inactivity timer.
- When the flow has reached a certain age, defined by a flow termination timer.

A flow record for path is created when a packet on a path is detected and no flow entry for that PPR-ID exists. I.e., flow records are not created and maintained for every PPR-ID that is known to the Network Element, only to PPR-IDs that are "active".

The following new IPFIX Information Elements need to be added:

- pprid-ipv4, with ipv4address as abstract data type.
- pprid-ipv6, with ipv6address as abstract data type.
- pprid-mpls, with unsigned32 as abstract data type.
- pprid-srv6, which is for further study.

Each of those fields need to be able to be supported as a flow key field.

The following statistics, as represented through corresponding Information Elements, will be of particular interest to export as part of IPFIX records with a PPR-ID as flow key:

- IE 1: octetDeltaCount - the number of octets on this path
- IE 2: packetDeltaCount - the number of incoming packets on this path
- IE 3: deltaFlowCount - the number of original flows routed via this path
- IE 21: flowEndSysUpTime - the relative time stamp of the last packet of the path flow (i.e. end of the observation period that is covered by the record)

- IE 22: flowStartSysUpTime - the relative time stamp of the first packet of the path flow (i.e. beginning of the observation period that is covered by the record)

- IE 132: droppedOctetDeltaCount - the number of octets of the path flow that were dropped by the node (during the observation period covered by the record).

- IE 133: droppedPacketDeltaCount - the number of packets of the path flow that were dropped by the node (during the observation period covered by the record).

Because a path's packets traverse every hop on the path, the observed path statistics are expected to generally be the same on every node across the path. Network providers may therefore choose to configure IPFIX path information export only on certain routers, for example at the network edge. That said, in certain circumstances it may make sense to export statistics from multiple nodes on a path and compare records for any discrepancies in order to diagnose or isolate operational anomalies (such as occurrence of packet loss).

7. A YANG Data Model for PPR Monitoring

7.1. Motivation and Overview

In addition to exporting path flow statistics, network providers need to be able to retrieve operational information about PPR paths as a whole. This includes information about when a path was created, how it came into being, and statistics maintained over the entire lifetime of the path (i.e. since its creation). For this purpose, a YANG Data Model for PPR Monitoring is defined, "ppr-statistics". The model is depicted in the following tree diagram. The tree diagram follows the notation defined in [RFC8340]

```
module: ietf-ppr-statistics
   +--ro ppr-stats
      +--ro num-pprs?   uint32
      +--ro ppr* [ppr-id]
         +--ro ppr-id                 ppr-id
         +--ro ppr-creation-time?     yang:date-and-time
         +--ro loose-or-strict?       ppr-path-type
         +--ro ppr-origin?            ppr-origin-proto
         +--ro ppr-packets?           uint64
         +--ro ppr-dropped-packets?   uint64
         +--ro ppr-octets?            uint64
         +--ro ppr-active-flows?      uint32
```

Figure 1: Tree diagram for establish-subscription-datastore-error-info

The data model contains the following:

- num-pprs: indicates the number of currently active PPRs on the node.

- ppr: the list of PPRs configured on the node, indexed by ppr-id, with the following entries:

  - ppr-creation-time: indicates the time the PPR came into being.

  - loose-or-strict: indicates whether the PPR is loose or strict.

  - ppr-origin: indicates the way in which the PPR came into being, i.e. through which method or protocol it was created.

  - ppr-packets: the number of packets that have forwarded on that path. (Wrapping around to 0 when the maximum is reached, per modulo semantics).

  - ppr-dropped packets: the number of packets on that path that have been dropped by this node. (Wrapping around to 0 when the maximum is reached, per modulo semantics).

  - ppr-octets: the number of octets that have been forwarded on that path. (Wrapping aound to 0 when the maxium is reached, per modulo semantics).

  - ppr-active-flows: The number of distinct flows that are currently active on that path.

7.2. YANG Data Model

```
<CODE BEGINS> file "ietf-ppr-statistics@2019-07-08.yang"
  module ietf-ppr-statistics {

  yang-version 1.1;
  namespace "urn:ietf:params:xml:ns:yang:ietf-ppr-statistics";

  prefix pprs;

  import ietf-yang-types {
    prefix yang;
  }

  import ietf-inet-types {
    prefix inet;
  }

  organization "IETF";
  contact
    "WG Web:   <http://tools.ietf.org/wg/rtgwg/>
     WG List:  <mailto:rtgwg@ietf.org>

     Author: Alexander Clemm
             <mailto:ludwig@clemm.org>

     Author: Padma Pillay-Esnault
             <mailto:padma@futurewei.com>

     Author: Uma Chunduri
             <mailto:uchundur@futurewei.com>";

  description
    "The YANG data model defines a set of statistics to be used for
     managing PPR.";

  revision 2019-07-08 {
    description
      "Initial revision";
    reference
      "RFC XXXX: Preferred Path Routing (PPR) OAM and Accounting";
  }

  typedef ppr-id {
    type union {
      type inet:ipv4-address;
      type inet:ipv6-address;
      type string {
```

```
          length "4..32";
        }
      }
      description
        "Identifies a PPR.  Depending on the type of PPR, a different
        format is used.";
    }

    typedef ppr-origin-proto {
      type string;
      description
        "Identifies the source of the PPR, i.e. how the PPR came into
        being.  Different values are TBD and the type itself is
        subject to change.";
    }

    typedef ppr-path-type {
      type enumeration {
        enum loose {
          description "Path type is loose";
        }
        enum strict {
          description "Path type is strict";
        }
      }
      description
        "The type of PPR path - loose or strict.";
    }

    container ppr-stats {
      config false;
      description
        "Top-level container for PPR statistics.";
      leaf num-pprs {
        type uint32;
        description
          "The number of currently active PPRs on the node.";
      }
      list ppr {
        key "ppr-id";
        description
          "The list of currently active PPRs on the node.";
        leaf ppr-id {
          type ppr-id;
          description
            "The identifier of the PPR.";
        }
        leaf ppr-creation-time {
```

```
          type yang:date-and-time;
          description
            "The precise time at which the PPR was created on the
            node.";
        }
        leaf loose-or-strict {
          type ppr-path-type;
          description
            "An indication whether the PPR is loose or strict.";
        }
        leaf ppr-origin {
          type ppr-origin-proto;
          description
            "The way in which the PPR came into being, i.e. through
            which method or protocol it was created.";
        }
        leaf ppr-packets {
          type uint64;
          description
            "The number of packets that have forwarded on that path.
            (Modulo semantics apply, i.e. the value of the leaf wraps
            around to 0 when the maximum uint64 is reached.)";
        }
        leaf ppr-dropped-packets {
          type uint64;
          description
            "The number of packets on that path that have been dropped
            by this node.  (Modulo semantics apply, i.e. the value of
            the leaf wraps around to 0 when the maximum uint64 is
            reached.)";
        }
        leaf ppr-octets {
          type uint64;
          description
            "The number of octets that have been forwarded on that
            path. (Modulo semantics apply, i.e. the value of the
            leaf wraps around to 0 when the maximum uint64 is
            reached.)";
        }
        leaf ppr-active-flows {
          type uint32;
          description
            "The number of distinct flows that are currently active
            on that path.";
        }
      }
    }
  }
```

<CODE ENDS>

8. Acknowledgements

tbd

9. IANA Considerations

9.1. IGP Path Attributes

This document requests the following new code points in IANA PPR Attributes TLV code-point registry for IS-IS, OSPFv2 and OSPFv3 protocols. IS-IS PPR Attributes are defined in [I-D.chunduri-lsr-isis-preferred-path-routing] and OSPF PPR attributes are defined in [I-D.chunduri-lsr-isis-preferred-path-routing].

Sub-TLV #	Sub-TLV Name
10	Packet Traffic Accounting (Section 5)
11	Traffic Statistics (Section 5)

add capability for loose traceroute support.

9.2. IPFIX Information Elements

IANA is requested to add the following entries to the IPFIX Information Elements Registry:

- pprid-ipv4, with ipv4address as abstract data type.
- pprid-ipv6, with ipv6address as abstract data type.
- pprid-mpls, with unsigned32 as abstract data type.
- pprid-srv6, whose abstract data type is for further study.

9.3. YANG Data Model

This document registers the following namespace URI in the "IETF XML Registry" [RFC3688]:

URI: urn:ietf:params:xml:ns:yang:ietf-ppr-statistics
Registrant Contact: The IESG.
XML: N/A; the requested URI is an XML namespace.

This document registers the following YANG module in the "YANG Module Names" registry [RFC6020]:

Name: ietf-ppr-statistics
Namespace: urn:ietf:params:xml:ns:yang:ietf-ppr-statistics
Prefix: pprs
Reference: draft-cpc-rtgwg-ppr-oam-XX.txt (RFC form)

10. Security Considerations

10.1. Path Trace and Ping

The ability to perform path traces and pings could be used by an attacker to discover details of a network. In addition, excessive amounts of traces and pings could be used by an attacker to try and exhaust network resources. Network providers therefore need to secure the ability to invoke trace and ping operations, requiring proper auhorization and authentication. Likewise, trace or ping requests originating from an untrusted source from outside the network edge should be dropped at the ingress edge.

10.2. IGPs and IPFIX

Security concerns for IS-IS are addressed in [RFC5304] and [RFC5310]. Further security analysis for IS-IS protocol is done in [RFC7645] with detailed analysis of various security threats and why [RFC5304] should not be used in the deployments. Advertisement of the additional information defined in this document introduces no new security concerns in IS-IS protocol. However as this extension is related to SR-MPLS and SRH data planes as defined in [I-D.ietf-spring-segment-routing], those particular data plane security considerations does apply here.

Existing security extensions for OSPF protocol are described in [RFC2328] and [RFC7684] apply to the extensions specified in this document. While OSPF is under a single administrative domain, there can be deployments where potential attackers have access to one or more networks in the OSPF routing domain. In these deployments, stronger authentication mechanisms such as those specified in [RFC7474] SHOULD be used.

This document also describes IPFX extensions that allow it to be used as a mechanism for IP Path Information Export. IPFIX security considerations apply, which are detailed in [RFC7011] section 11.

10.3. YANG Data Model

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 Network Configuration Access Control Model (NACM) [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.

Some of the readable data nodes in this YANG module may be considered sensitive or vulnerable in some network environments. It is thus important to control read access (e.g., via get, get-config, or notification) to these data nodes. The readable data nodes are defined under the container "ppr-stats". The sensitivity/ vulnerability concerns providing an unauthorized attacker with internals about the network, specifically exposure of PPR IDs that are installed on a network node and insight about the network traffic that occurs over these nodes.

11. References

11.1. Normative References

[I-D.chunduri-lsr-isis-preferred-path-routing]
Chunduri, U., Li, R., White, R., Tantsura, J., Contreras, L., and Y. Qu, "Preferred Path Routing (PPR) in IS-IS", draft-chunduri-lsr-isis-preferred-path-routing-03 (work in progress), May 2019.

[I-D.chunduri-lsr-ospf-preferred-path-routing]
Chunduri, U., Qu, Y., White, R., Tantsura, J., and L. Contreras, "Preferred Path Routing (PPR) in OSPF", draft-chunduri-lsr-ospf-preferred-path-routing-03 (work in progress), May 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>.

[RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, DOI 10.17487/RFC3688, January 2004, <https://www.rfc-editor.org/info/rfc3688>.

[RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <https://www.rfc-editor.org/info/rfc4443>.

[RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, DOI 10.17487/RFC6020, October 2010, <https://www.rfc-editor.org/info/rfc6020>.

[RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <https://www.rfc-editor.org/info/rfc7011>.

[RFC8029] Kompella, K., Swallow, G., Pignataro, C., Ed., Kumar, N., Aldrin, S., and M. Chen, "Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures", RFC 8029, DOI 10.17487/RFC8029, March 2017, <https://www.rfc-editor.org/info/rfc8029>.

[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>.

[RFC8287] Kumar, N., Ed., Pignataro, C., Ed., Swallow, G., Akiya, N., Kini, S., and M. Chen, "Label Switched Path (LSP) Ping/Traceroute for Segment Routing (SR) IGP-Prefix and IGP-Adjacency Segment Identifiers (SIDs) with MPLS Data Planes", RFC 8287, DOI 10.17487/RFC8287, December 2017, <https://www.rfc-editor.org/info/rfc8287>.

[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>.

[RFC8341] Bierman, A. and M. Bjorklund, "Network Configuration Access Control Model", STD 91, RFC 8341, DOI 10.17487/RFC8341, March 2018, <https://www.rfc-editor.org/info/rfc8341>.

11.2. Informative References

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <https://www.rfc-editor.org/info/rfc2328>.

[RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, DOI 10.17487/RFC5304, October 2008, <https://www.rfc-editor.org/info/rfc5304>.

[RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, DOI 10.17487/RFC5310, February 2009, <https://www.rfc-editor.org/info/rfc5310>.

[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>.

[RFC7474] Bhatia, M., Hartman, S., Zhang, D., and A. Lindem, Ed., "Security Extension for OSPFv2 When Using Manual Key Management", RFC 7474, DOI 10.17487/RFC7474, April 2015, <https://www.rfc-editor.org/info/rfc7474>.

[RFC7645] Chunduri, U., Tian, A., and W. Lu, "The Keying and Authentication for Routing Protocol (KARP) IS-IS Security Analysis", RFC 7645, DOI 10.17487/RFC7645, September 2015, <https://www.rfc-editor.org/info/rfc7645>.

[RFC7684] Psenak, P., Gredler, H., Shakir, R., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute Advertisement", RFC 7684, DOI 10.17487/RFC7684, November 2015, <https://www.rfc-editor.org/info/rfc7684>.

[RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <https://www.rfc-editor.org/info/rfc8040>.

[RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <https://www.rfc-editor.org/info/rfc8446>.

Authors' Addresses

Alexander Clemm
Futurewei
2330 Central Expressway
Santa Clara, CA 95050
USA

Email: ludwig@clemm.org

Padma Pillay-Esnault
Futurewei
2330 Central Expressway
Santa Clara, CA 95050
USA

Email: padma@futurewei.com

Uma Chunduri
Futurewei
2330 Central Expressway
Santa Clara, CA 95050
USA

Email: umac.ietf@gmail.com