Interdomain Routing Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 24, 2020

M. Liu
Y. Wang
Huawei
R. Pang
China Unicom

March 23, 2020

BGP Flow Specification Extensions to Enable In-situ Flow Information Telemetry (IFIT) draft-liu-idr-flowspec-ifit-04

Abstract

BGP Flowspec mechanism propogates both traffic Flow Specifications and Traffic Filtering Actions by making use of the Border Gateway Protocol Network Layer Reachability Information (BGP NLRI) and the BGP Extended Community encoding formats. In order to address the automatic deployment of IPv4 unicast and VPNv4 unicast on-path flow telemetry as well as IPv6 families, this document specifies a new BGP Extended Community named IFIT Action Specific Extended Community to distribute In-situ Flow Information Telemetry (IFIT) actions.

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 RFC 2119 [RFC2119].

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 24, 2020.

Table of Contents

1. Introduction

At present, a family of on-path flow telemetry techniques referred in [I-D.song-opsawg-ifit-framework] are emerging, including In-situ OAM (IOAM) [I-D.ietf-ippm-ioam-data], Postcard-Based Telemetry (PBT) [I-D.song-ippm-postcard-based-telemetry], IOAM Direct Export (DEX) [I-D.ioamteam-ippm-ioam-direct-export], Enhanced Alternate Marking (EAM) [I-D.zhou-ippm-enhanced-alternate-marking]. we categorize these on-path telemetry echniques as the hybrid OAM type I according to the classification defined in [RFC7799]. These techniques provide flow information on the entire forwarding path on a per-packet basis in real time, which are useful for application-aware network operations not only in data center and enterprise networks but also in carrier networks which may cross multiple domains. The data provided by on-

path telemetry are especially useful for network operations in aspects of SLA compliance, service path enforcement, fault diagnosis, and network resource optimization, etc. In IFIT reflection-loop architecture [I-D.song-opsawg-ifit-framework], an IFIT functionality needs to choose a suite of telemetry tecchniques and enable initial techniques to the data plane in accordance to the monitoring and measurement requirements. Then the IFIT head nodes need to determine the target flows and packets to apply the IFIT-specific functions and the telemetry data sets.

However, enabling only a single underlying on-path telemetry technique may lead to defective result. A comprehensive solution needs the flexibility to switch between different underlying techniques and enable different IFIT option types to adapt to different network conditions and different application requirements. It's useful for application-aware network operations to enable desired IFIT option types to the target flows dynamically.

As we know, Dissemination of Flow Specification Rules [I-D.ietf-idr-rfc5575bis] provides a protocol extension for propagation of traffic flow information for the purpose of rate limiting, filtering, shaping, classifying or redirecting. And BGP extended community encoding formats can be used to propagate traffic filtering actions along with the flow specification NLRI. Those traffic filtering actions encode actions a routing system can take if the packet matches the flow specifications. And the other document [I-D.ietf-idr-flow-spec-v6] extends BGP Flowspec [I-D.ietf-idr-rfc5575bis] and to make it also usable and applicable to IPv6 data packets.

From an operational perspective, the utilization of BGP Flowspec as the carrier for the specific flow information allows a network service provider to reuse BGP route distribute infrastructure. Therefore, this document defines the IFIT Action Specific Extended Community to enable the application of IPv4 unicast and VPNv4 unicast on-path flow telemetry as well as IPv6 families.

2. Terminologies

IFIT: In-situ Flow Information Telemetry

NLRI: Network Layer Reachability Information

3. Motivation

The IFIT functionality, which enables the future autonomous network operation, will pick one of proper In-situ Flow Information Telemetry techniques and apply a flow, packet, and data selection policy to

monitor the specific traffic flow for application-aware network operation. In current deployments, there have been relatively static methods, ACL-like CLI and NETCONF with YANG model to configure the specific flows or packets to be monitored on the relevant IFIT-capable nodes. However, with the evolution of Intent-based and autonomous network operation, and the trends of network virtualization, network convergence, and packet-optical integration, the future data plane telemetry will support an on-demand and interactive fashion. Flexibility and extensibility of telemetry data defining and acquiring must be considered. Therefore, flexible deployment of IFIT option types based on the real-time telemetry data analysis results and telemetry requirements of different applications is needed.

BGP Flowspec mechanism is preferred in the reflective-loop network telemetry system. This document defines IFIT Action Specific Extended Community to enable IFIT functionality for the relevant flows that match the Traffic Flow Specifications along with the BGP NLRI defined in [I-D.ietf-idr-rfc5575bis] and [I-D.ietf-idr-flow-spec-v6]. The IFIT Action Specific Extended Community instructs a routing system to add the IFIT-Option-Types into packets of flows and update relevant IFIT-Data-Fields in packets that traverse.

4. IFIT Action Specific Extended Community

This section defines a new BGP Extended Community and different sub-types in accordance with different IFIT option types.

Traffic Filtering Actions that are standardized as BGP Extended Community, which is encoded as an 8-octet quantity containing Type field and Value field [RFC4360]. The Types are to be assigned by IANA registry. The Value field contains Traffic Filtering Action values.

In the IFIT framework architecture, there are a few of available option types for the specified traffic flow, e.g. IOAM pre-allocated/incremental trace [I-D.ietf-ippm-ioam-data], IOM Edge-to-Edge [I-D.ietf-ippm-ioam-data], IOAM Direct Export (DEX) [I-D.ioamteam-ippm-ioam-direct-export], Enhanced Alternate Marking (EAM) [I-D.zhou-ippm-enhanced-alternate-marking], etc. As different IFIT option types have different formats of parameters, following defines the Type and various sub-types of Extended Communities in accordance with different IFIT option types.

- Type tt: IFIT Action

- Sub-type ss1: IOAM Pre-allocated Trace Option

- Sub-type ss2: IOAM Incremental Trace Option
- Sub-type ss3: IOAM DEX Option
- Sub-type ss4: IOAM Edge-to-Edge Option
- Sub-type ss5: Enhanced Alternate Marking Option

IFIT Action do not interfere with other BGP Flow Specification Traffic Filtering Action defined in [I-D.ietf-idr-rfc5575bis] and [I-D.ietf-idr-flow-spec-v6].

In the following sections, the different IFIT Action Specific Extened Communities encoding formats are presented.

4.1. IOAM Pre-allocated Trace Option sub-type

The IOAM tracing data is expected to be collected at every node that a packet traverses to ensure visibility into the entire path a packet takes within an IOAM domain. The pre-allocated tracing option will create pre-allocated space for each node to populate its information.

The format of IOAM pre-allocated trace option Extended Community is defined as follows:

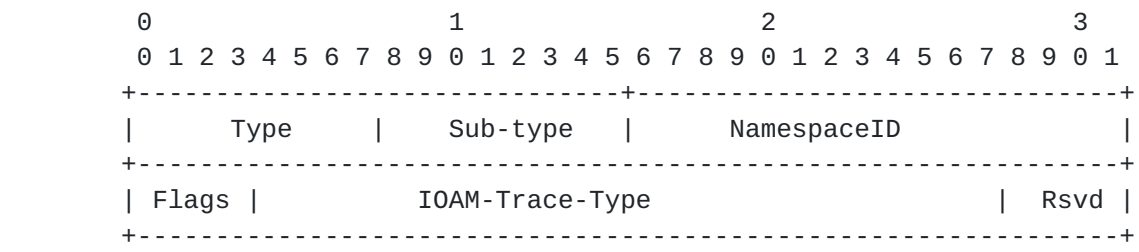

Fig. 1 IOAM Pre-allocated Trace Option Extended Community Encoding

Where:

Type: to be assigned by IANA.

Sub-type: to be assigned by IANA.

Namespace ID: A 16-bit identifier of an IOAM-Namespace. The definition is the same as described in section 4.4 of[I-D.ietf-ippm-ioam-data] .

Flags: A 4-bit field. The definition is the same as described in [I-D.ietf-ippm-ioam-flags] and section 4.4 of [I-D.ietf-ippm-ioam-data].

IOAM-Trace-Type: A 24-bit identifier which specifies which data types are used in the node data list. The definition is the same as described in section 4.4 of [I-D.ietf-ippm-ioam-data].

Rsvd: A 4-bit field reserved for further usage. It should be set to zero and must be ignored during decoding.

4.2. IOAM Incremental Trace Option sub-type

The incremental tracing option contains a variable node data fields where each node allocates and pushes its node data immediately following the option header.

The format of IOAM incremental trace option Extended Community is defined as follows:

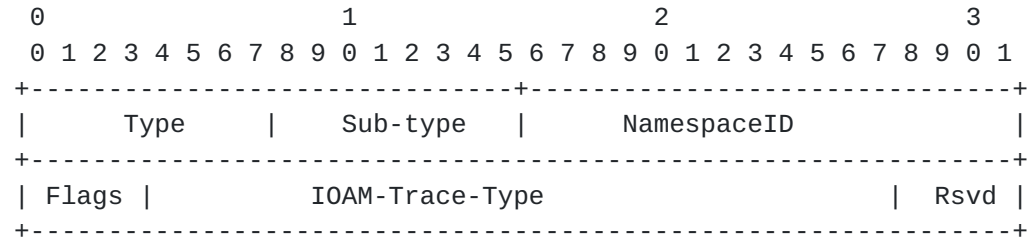

Fig. 2 IOAM Incremental Trace Option Extended Community Encoding

Where:

Type: to be assigned by IANA.

Sub-type: to be assigned by IANA.

All the other fields definistion is the same as the pre-allocated trace option Extended Community in section 3.2.1.

4.3. IOAM DEX Option sub-type

The DEX option is used as a trigger to export IOAM data to a collector. Moreover, IOAM nodes may send exported data for all traversing packets that carry the DEX option, or may selectively export data only for a subset of these packets. The DEX option specifies which data fields should be exported to the collector, as specified in Section 3.2 of [I-D.ioamteam-ippm-ioam-direct-export].

The format of IOAM DEX option Extended Community is defined as follows:

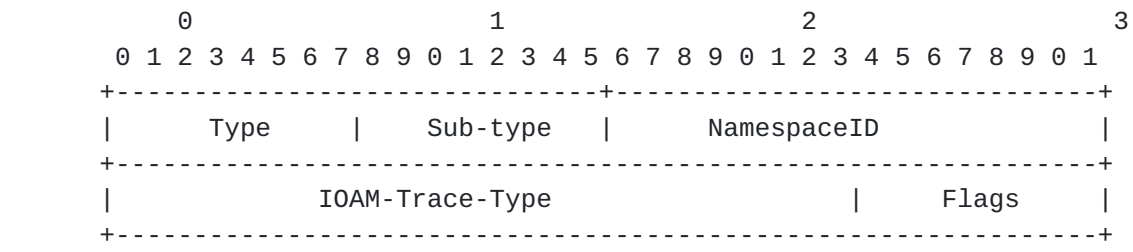

Fig. 3 IOAM DEX Option Extended Community Encoding

Where:

Type: to be assigned by IANA.

Sub-type: to be assigned by IANA.

Namespace-ID: a 16-bit identifier of the IOAM namespace, as defined in section 4.4 of [I-D.ietf-ippm-ioam-data].

IOAM-Trace-Type: a 24-bit identifier which specifies which data fields should be exported. The format of this field is as defined in section 4.4 of[I-D.ietf-ippm-ioam-data].

Flags: A 8-bit field, comprised of 8 one-bit subfields. Flags are allocated by IANA.

4.4. IOAM Edge-to-Edge Option sub-type

The IOAM edge to edge option is to carry data that is added by the IOAM encapsulating node and interpreted by IOAM decapsulating node.

The format of IOAM edge-to-edge option Extended Community is defined as follows:

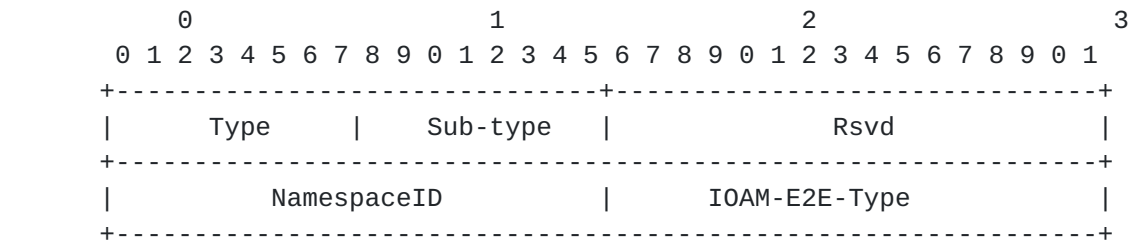

Fig. 4 IOAM Edge-to-Edge Option Extended Community Encoding

Where:

Type: to be assigned by IANA.

Sub-type: to be assigned by IANA.

Namespace ID: A 16-bit identifier of an IOAM-Namespace. The definition is the same as described in section 4.6 of[I-D.ietf-ippm-ioam-data].

IOAM-E2E-Type: A 16-bit identifier which specifies which data types are used in the E2E option data. The definition is the same as described in section 4.6 of [I-D.ietf-ippm-ioam-data].

Rsvd: A 16-bit field reserved for further usage. It should be set to zero.

4.5. Enhanced Alternate Marking Option sub-type

The Alternate Marking [RFC8321] technique is an hybrid performance measurement method and can be used to measure packet loss, latency, and jitter on live traffic because it is based on marking consecutive batches of packets.

The Enhanced Alternate Marking (EAM) [I-D.zhou-ippm-enhanced-alternate-marking] defines data fields for the alternate marking with enough space, in particular for Postcard-based Telemetry. More information can be considered within the alternate marking field to facilitate the efficiency and ease the deployment.

The format of EAM Option Extended Community is defined as follows:

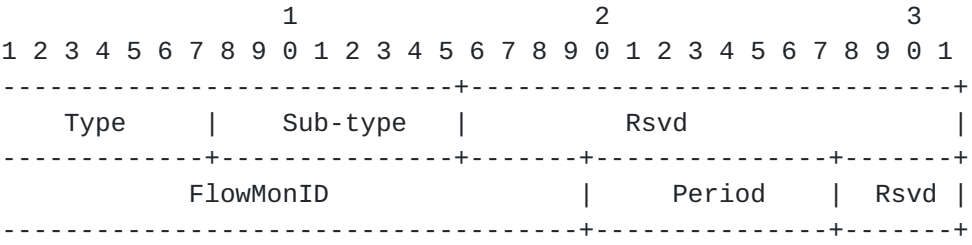

Fig. 5 Enhanced Alternate Marking Option Extended Community Encoding

Where:

Type: to be assigned by IANA.

Sub-type: to be assigned by IANA.

FlowMonID: A 20-bit identifier to uniquely identify a monitored flow within the measurement domain. The definition is the same as described in section 2 of [I-D.zhou-ippm-enhanced-alternate-marking].

Period: A 8-bit field. Time interval between two alternate marking period. The unit is second.

Rsvd: A 4-bit field reserved for further usage. It should be set to zero.

5. IANA Considerations

This document requests a new Transitive Extended Community Type and five new registery sub-types. The new Transitive Extended Community Type name shall be "IFIT Action Extended Community (Sub-Types are defined in the "IFIT Action Extended Community Sub-Type" registery)".

Type Value	Name
TBD	IFIT Action

Sub-type Value	Name
TBD	IOAM Pre-allocated Trace Option
TBD	IOAM Incremental Trace Option
TBD	IOAM DEX Option
TBD	IOAM E2E Option
TBD	Enhanced Alternate Marking

6. Security Considerations

No new security issues are introduced to the BGP Flow Specifications and Traffic Filtering Action in [I-D.ietf-idr-flow-spec-v6] and [I-D.ietf-idr-rfc5575bis].

7. Acknowledgements

TBD.

8. References

8.1. Normative References

[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>.

[RFC4360] "BGP Extended Communities Attribute", <https://www.rfc-editor.org/info/rfc4360>.

[RFC7799] "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", <https://www.rfc-editor.org/info/rfc7799>.

[RFC8321] "Alternate-Marking Method for Passive and Hybrid Performance Monitoring", <https://www.rfc-editor.org/info/rfc8321>.

8.2. Informative References

[I-D.ietf-idr-flow-spec-v6] "Dissemination of Flow Specification Rules for IPv6", <https://datatracker.ietf.org/doc/draft-ietf-idr-flow-spec-v6/>.

[I-D.ietf-idr-rfc5575bis] "Dissemination of Flow Specification Rules", <https://datatracker.ietf.org/doc/draft-ietf-idr-rfc5575bis/>.

[I-D.ietf-ippm-ioam-data] "Data Fields for In-situ OAM", <https://datatracker.ietf.org/doc/draft-ietf-ippm-ioam-data/>.

[I-D.ioamteam-ippm-ioam-direct-export] "In-situ OAM Direct Exporting", <https://datatracker.ietf.org/doc/draft-ioamteam-ippm-ioam-direct-export/>.

[I-D.song-ippm-postcard-based-telemetry] "Postcard-based On-Path Flow Data Telemetry", <https://datatracker.ietf.org/doc/draft-song-ippm-postcard-based-telemetry/>.

[I-D.song-opsawg-ifit-framework] "In-situ Flow Information Telemetry Framework", <https://datatracker.ietf.org/doc/draft-song-opsawg-ifit-framework/>.

[I-D.zhou-ippm-enhanced-alternate-marking]
"Enhanced Alternate Marking Method",
<https://datatracker.ietf.org/doc/draft-zhou-ippm-enhanced-alternate-marking/>.

Authors' Addresses

Min Liu
Huawei
156 Beiqing Rd., Haidian District
Beijing
China

Email: lucy_liumin@huawei.com

Yali Wang
Huawei
156 Beiqing Rd., Haidian District
Beijing
China

Email: wangyali11@huawei.com

Ran Pang
China Unicom
9 Shouti South Rd., Haidian District
Beijing
China

Email: pangran@chinaunicom.cn