IETF
Internet-Draft
Intended status: Informational
Expires: 23 April 2026

Y. Cui
Tsinghua University
Y. Gao
L. Zhang
Zhongguancun Laboratory

20 October 2025

BGP Flow Specification Extension for Feedback Binding
draft-cui-idr-flowspec-feedback-binding-00

Abstract

This document specifies a BGP Flow Specification extension that conveys per-route feedback binding for a FlowSpec route using the BGP Extended Community attribute. The proposed mechanism introduces a single Feedback Action, encoded as a Generic Transitive Extended Community, which enables downstream routers to report telemetry information or operational events associated with a FlowSpec rule. The Feedback Action carries parameters including a Feedback Identifier (FID), a window exponent (WINC) that defines the periodic aggregation interval, an event flag, and a scope selector to control where feedback is generated. These parameters are attached to the FlowSpec route and are propagated across AS boundaries unchanged. This document focuses on the signaling aspect; a companion document may define how feedback information is exported as part of a network telemetry framework (e.g., leveraging the BGP Monitoring Protocol (BMP)) or equivalent mechanisms to report periodic and event-driven feedback keyed by the FID.

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 23 April 2026.

Table of Contents

1. Introduction

BGP Flow Specification (FlowSpec) defines a method for distributing traffic filtering rules using BGP, enabling operators to deploy network-wide traffic management and DDoS mitigation policies in a scalable manner.
The existing versions, FlowSpec (FSv1) RFC8955 [RFC8956] and FlowSpec Version 2 (FSv2) [draft-ietf-idr-flowspec-v2-04], allow the advertisement of flow-matching conditions and actions to drop, rate-limit, or redirect traffic.
These mechanisms are widely used for dynamic DDoS mitigation and policy enforcement across Autonomous Systems (ASes).

However, current FlowSpec deployments lack a standardized mechanism for feedback reporting that allows the originator of a rule to obtain operational visibility—such as installation status, traffic hit counts, or rule effectiveness—from downstream routers. Without such feedback, operators must rely on out-of-band telemetry or vendor-specific mechanisms, leading to fragmented monitoring and difficulty in verifying the success of mitigation actions, especially in multi-domain environments.

To address this limitation, this document introduces a new Feedback Action that extends the FlowSpec capability to include feedback control and telemetry binding at the route level. By signaling feedback parameters directly within BGP, the originator can request periodic or event-driven reports about the operational state of specific FlowSpec rules. This enhancement enables a closed-loop control paradigm where policy dissemination and enforcement can be continuously monitored and optimized.

This document focuses solely on the signaling aspect of feedback within BGP FlowSpec. The actual transport of feedback data—such as telemetry reports or event notifications—is out of scope and may be realized using the BGP Monitoring Protocol (BMP) [RFC7854] or other telemetry frameworks compliant with the Network Telemetry Framework.

2. Terminology

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. Definitions and Acronyms

* DDoS: Distributed Denial of Service

* NLRI: Network Layer Reachability Information

* FSv1: Flow Specification Version 1, defined in [RFC8955] and [RFC8956]

* FSv2: Flow Specification Version 2 define in [draft-ietf-idr-flowspec-v2-04]

* Telemetry: A framework for real-time or streaming collection of network operational data, as defined in [RFC9232].

* BMP: BGP Monitoring Protocol [RFC7854], a telemetry protocol used for exporting BGP operational and statistical data.

4. Overview

This specification defines the Feedback Action, a new BGP FlowSpec action encoded as a Generic Transitive Extended Community (Type 0x80).
The Feedback Action allows the originator of a FlowSpec route to convey parameters that instruct downstream routers how and where to generate telemetry feedback for that route.

The Feedback Action carries four parameters compactly encoded in its 6-octet Value field:

* Feedback Identifier (FID) A unique identifier that associates telemetry reports with a specific FlowSpec rule.

* Window Exponent (WINC) Defines the aggregation interval as 2^WINC seconds for periodic reporting.

* Event Flag (E) Controls whether feedback is periodic (00) or event-only (01).

* Scope (S) Specifies the feedback domain: Global, Inter-AS, or Intra-AS.

When attached to a FlowSpec NLRI, the Feedback Action expresses the originator’ s intent for feedback generation.
It is transitive, ensuring that all capable routers along the propagation path can interpret and act upon the feedback instruction, while non-supporting routers transparently forward it unchanged to maintain compatibility.

This signaling mechanism does not define the feedback transport itself.
Actual telemetry or event reports may be exported through BMP or other protocols aligned with [RFC9232], allowing integration into existing network telemetry infrastructures without altering FlowSpec semantics.

In summary, the Feedback Action augments FlowSpec with a lightweight, interoperable feedback control mechanism that enables closed-loop telemetry, enhances operational visibility, and supports both single-domain and multi-domain DDoS defense deployments with minimal protocol overhead. This extension is suitable forboth FSv1 and FSv2.

5. Feedback Action Encoding

The Feedback Action is conveyed using a single BGP Generic Transitive Extended Community, attached to the FlowSpec NLRI.
This action encodes a compact set of parameters that define the feedback reporting behavior for a specific FlowSpec route, including a Feedback Identifier (FID), a Window Exponent (WINC), an Event Flag (E), and a Scope Selector (S).
The action MUST include at least the FID, WIND, and Event flag for a valid binding; the Scope Selector is OPTIONAL. If multiple communities with the same tag are present, receivers SHOULD use the first one and ignore duplicates.

5.1. Encoding Format

The Feedback Action uses the standard IPv4 Extended Community encoding format, as illustrated below:

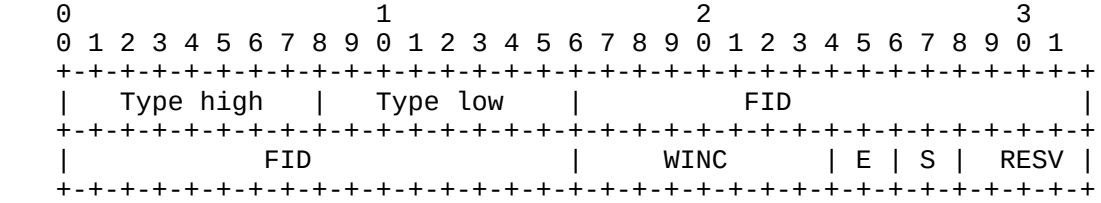

Figure 1: Feedback Action Encoding format

* Type high (1 octet): 0x80 — indicates a Generic Transitive Extended Community as defined in [RFC4360].

* Type low (1 octet): TBD (Feedback Action Sub-Type) — value to be assigned by IANA.

* FID (4 octets): A 32-bit Feedback Identifier, unique within the originating AS.
  It serves as the key for feedback correlation between the FlowSpec rule and its associated telemetry data.
  A value of zero is invalid and MUST be ignored.

* WINC (1 octet): Window Exponent, defining the periodic aggregation window as 2^WINC seconds.
  Valid range: 031 (corresponding to 1s to approximately 24 days).
  Values above 31 are reserved and MUST be ignored.

* E (2 bits): Event flag, specifying the feedback triggering mode:

  - 00 — Periodic feedback enabled.

- 01 — Event-only feedback (e.g., on rule install, remove, or error).
- 10 and 11 — Reserved for future use (MUST be set to zero on transmission and ignored on receipt).

* S (2 bits): Scope selector, specifying where feedback reports are generated:
  - 00 — Global (default; feedback from all capable receivers).
  - 01 — Inter-AS (only from downstream ASes).
  - 10 — Intra-AS (only within the same AS).
  - 11 — Reserved.
* *RESV (4 bits):* Reserved for future extensions.
  MUST be set to zero by the sender and ignored by the receiver.

5.2. Fields Descriptions

5.2.1. Feedback Identifier (FID)

The FID uniquely identifies the feedback stream corresponding to the FlowSpec route.
The combination of (Originator ASN, FID) forms a globally unique key for feedback correlation.
This identifier allows the originator to distinguish multiple FlowSpec rules and their respective telemetry results.
Routers that do not recognize this field MUST forward the Extended Community unchanged.

5.2.2. Window Length (WIND)

The WINC parameter determines the frequency of periodic feedback reports.
The interval is expressed as an exponent base 2, providing scalability from second to multi-day granularity.
Receivers SHOULD align their reporting schedule to the nearest integer multiple of 2^WINC seconds to maintain synchronization.

5.2.3. Event Flag (E)

The Event Flag specifies whether the feedback should be triggered periodically or only upon discrete events.
When E=01, routers generate reports only when specific control-plane or data-plane events occur,
such as rule installation, withdrawal, update failure, or error detection.
This allows efficient on-demand visibility without unnecessary periodic reporting overhead.

5.2.4. Scope Selector (S)

The Scope Selector defines the domain from which feedback reports are generated.
This enables fine-grained control over where telemetry is collected:
* Global (00): Feedback is expected from all capable routers across AS boundaries.
* Inter-AS (01): Feedback is limited to routers in downstream ASes.
* Intra-AS (10): Feedback is generated only within the same administrative domain.
If a receiver does not match the specified scope, it MAY silently ignore the Feedback Action.

5.2.5. Reserved Bits (RESV)

The RESV field (4 bits) is reserved for future extensions.
It MUST be set to zero on transmission and ignored on receipt.

5.3. Validation Rules

Receivers MUST validate the Feedback Action upon reception to ensure it conforms to the expected encoding and operational constraints.
The following validation rules apply:

* The Feedback Action MUST be attached to a valid FlowSpec NLRI.

* The FID field MUST be non-zero; a value of zero indicates an invalid binding and the attribute MUST be ignored.

* The WINC (Window Exponent) field MUST NOT exceed 31. Values above this limit are considered invalid and MUST be ignored.

* Reserved or undefined values in the E (Event Flag), S (Scope Selector), or RESV bits MUST be set to zero on transmission and MUST be ignored on receipt.

* The total length of the Extended Community attribute MUST conform to the standard 8-octet format defined in [RFC4360]. Any deviation from this format MUST cause the attribute to be ignored.

* Malformed attributes or decoding errors MUST NOT result in session reset.
  Such attributes SHOULD be logged locally for operational visibility and SHOULD be propagated unchanged to preserve transitivity.

* If multiple Feedback Actions are present for the same NLRI, only the first instance SHOULD be processed; subsequent duplicates SHOULD be ignored.

A Feedback Action that fails any of the above validation checks MUST be silently discarded and MUST NOT affect normal FlowSpec rule installation or propagation.

6. Propagation and Usage

The Feedback Action is attached to the BGP UPDATE message that carries the FlowSpec NLRI and MUST be propagated unchanged across all BGP peers, including across AS boundaries.
This ensures consistent feedback signaling while preserving the original semantics and reachability of the FlowSpec route.

Routers that support feedback reporting SHOULD generate telemetry or event reports according to the parameters conveyed in the attribute. When the local scope matches the value of _S_ and the Event Flag indicates periodic or event-driven reporting, feedback reports SHOULD be generated accordingly.
The absence of feedback reports MUST NOT be interpreted as an error, and the FlowSpec rule remains valid for traffic enforcement.

If the attribute cannot be parsed or fails validation, it MUST be silently ignored and MUST NOT trigger a BGP session reset.
Such attributes SHOULD be logged locally for operational visibility and MUST be propagated unchanged to preserve transitivity.
Routers MUST NOT alter or regenerate the Feedback Action during re-advertisement.

This mechanism provides a lightweight and interoperable means of achieving closed-loop telemetry for FlowSpec deployments, enabling standardized in-band feedback signaling while maintaining full backward compatibility with existing BGP implementations.

7. Security Considerations

This extension inherits the security properties of BGP FlowSpec and Large Communities. Potential issues include:

* Resource exhaustion: Receivers SHOULD rate-limit feedback generation and ignore excessive requests.

* Privacy: Feedback may reveal traffic patterns; the companion telemetry document SHOULD recommend encryption.

* Amplification: Malicious bindings could trigger unwanted reports; originators SHOULD authenticate telemetry receivers.

Operators SHOULD filter invalid or unauthorized communities at AS borders using ingress/egress policies.

8. IANA Considerations

This document requests IANA to assign a new Sub-Type called "feedback action" under the “BGP Extended Communities” registry:

Type	Sub-Type	Name	Reference
TBD	TBD	Feedback Action	This document

Table 1

9. Normative References

[RFC8955] Loibl, C., Hares, S., Raszuk, R., McPherson, D., and M. Bacher, "Dissemination of Flow Specification Rules", RFC 8955, DOI 10.17487/RFC8955, December 2020, <https://www.rfc-editor.org/rfc/rfc8955>.

[RFC8956] Loibl, C., Ed., Raszuk, R., Ed., and S. Hares, Ed., "Dissemination of Flow Specification Rules for IPv6", RFC 8956, DOI 10.17487/RFC8956, December 2020, <https://www.rfc-editor.org/rfc/rfc8956>.

[RFC9117] Uttaro, J., Alcaide, J., Filsfils, C., Smith, D., and P. Mohapatra, "Revised Validation Procedure for BGP Flow Specifications", RFC 9117, DOI 10.17487/RFC9117, August 2021, <https://www.rfc-editor.org/rfc/rfc9117>.

[RFC8092] Heitz, J., Ed., Snijders, J., Ed., Patel, K., Bagdonas, I., and N. Hilliard, "BGP Large Communities Attribute", RFC 8092, DOI 10.17487/RFC8092, February 2017, <https://www.rfc-editor.org/rfc/rfc8092>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/rfc/rfc8126>.

[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/rfc/rfc2119>.

[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/rfc/rfc8174>.

[draft-ietf-idr-flowspec-v2-04]
Hares, S., 3rd, D. E. E., Yadlapalli, C., and S. Maduschke, "BGP Flow Specification Version 2", April 2024, <https://datatracker.ietf.org/doc/draft-ietf-idr-flowspec-v2/04/>.

Authors' Addresses

Yong Cui
Tsinghua University
Beijing, 100084
China
Email: cuiyong@tsinghua.edu.cn
URI: http://www.cuiyong.net/

Yujia Gao
Zhongguancun Laboratory
Beijing, 100094
China
Phone: +86-185-1028-7458
Email: gaoyj@zgclab.edu.cn

Lei Zhang
Zhongguancun Laboratory
Beijing, 100094
China
Email: zhanglei@zgclab.edu.cn