Packet Content Filter for BGP FlowSpec

draft-cui-idr-content-filter-flowspec-04

Yong Cui, Tsinghua University

Yujia Gao, Zhongguancun Laboratory

Susan Hares, Hickory Hill Consulting

May 11th, 2026

Current Draft Status

- 2024/3 first proposed packet content filter @IETF 119
- 2024/5 join FSv2 design team work
- 2024/6 conduct software router function development and testbed validation
- 2024/7 introduce the latest work @IETF 120
- 2025/9 discuss with vendors and operators to optimize filter definitions
- 2024/11 introduce the work progress @IETF 121
- 2025/10 introduce the latest work @IETF 124
- 2026/2 update scaling, operational, security considerations @Interim meeting
- 2026/5 Update draft @Interim meeting

Problem Statement

- DDoS attack intensity has surged, with a record peak of 22.2 Tbps, far beyond scrubbing centers' defense capability
- Some real world attacks captured in operator networks show fixed payload (e.g., ACK flood, SSDP attack)

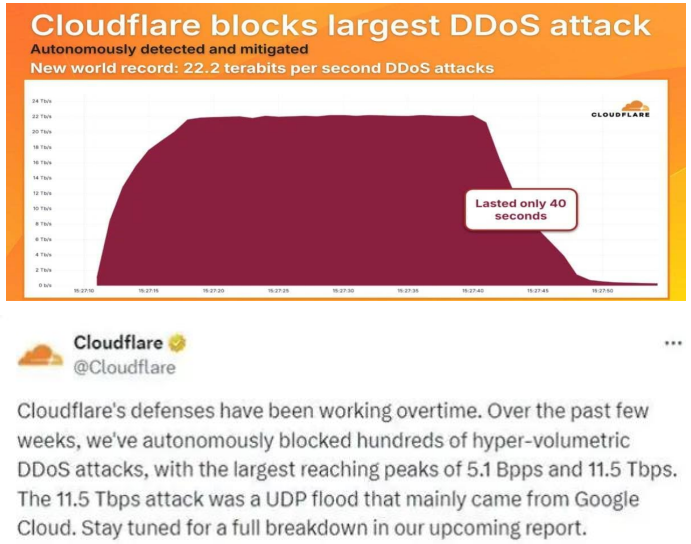

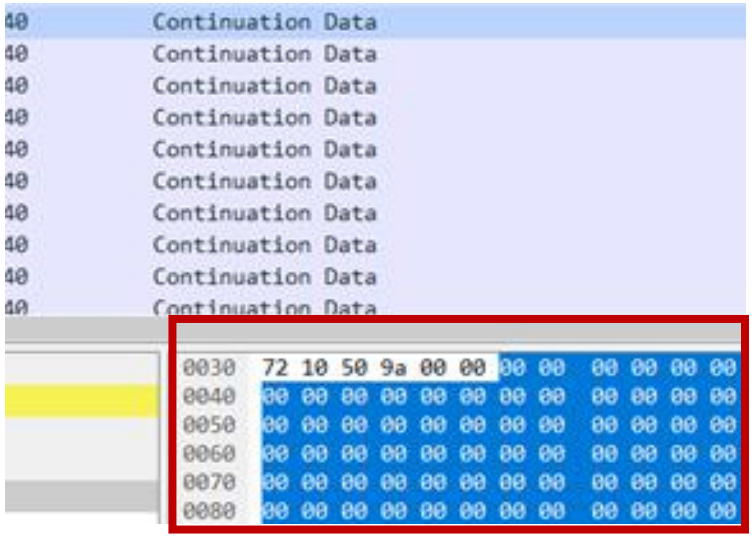

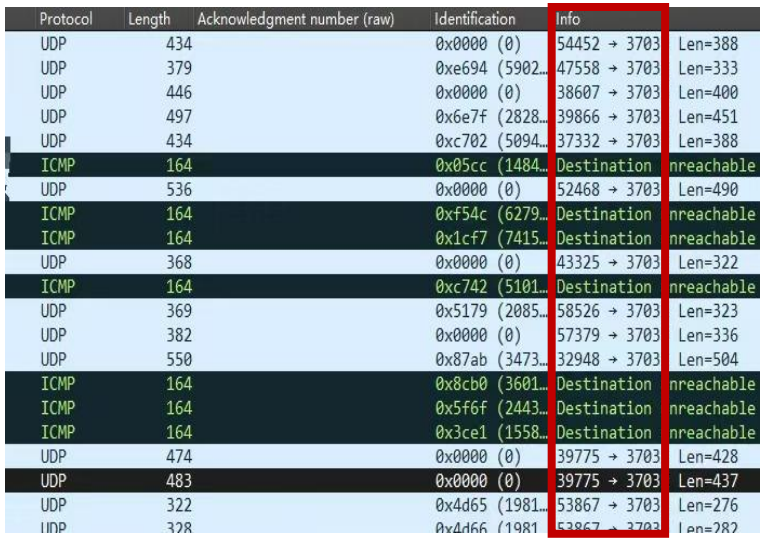

A new type of FlowSpec filter is needed to defend against DDoS attacks with specific payload characteristics.

Packet Content Filter

- NLRI Encoding:

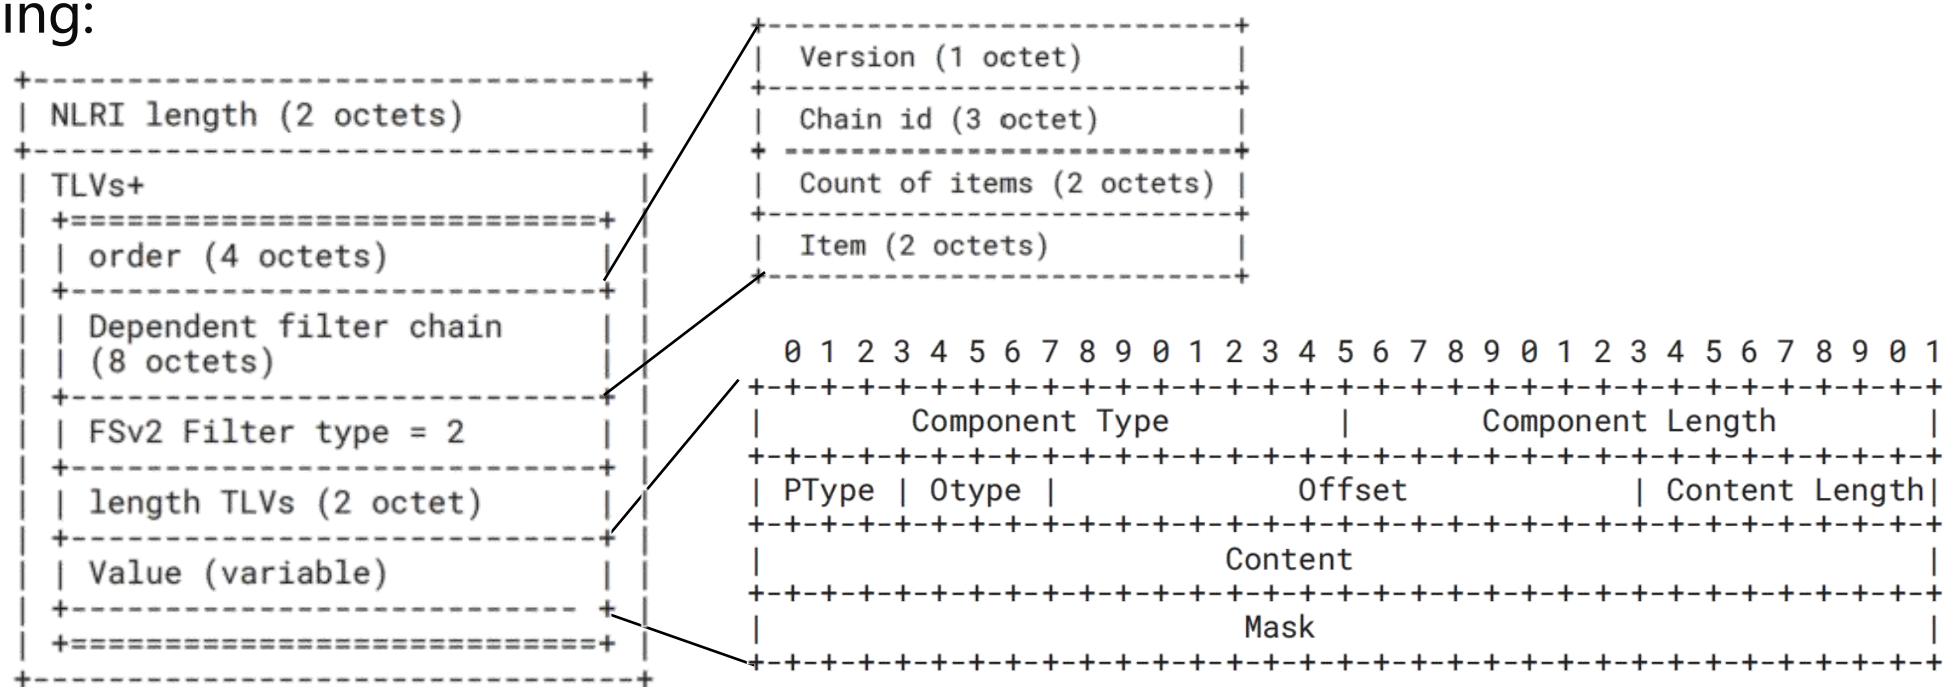

- Ordering rules:
  a) Filters with a specific user order number would be ordered by the user order
  b) Filters with same (or no) user order would be ordered by the default order:

Content-length(⬇)→Otype(⬇)→Offset (⬇) →Content(⬆)

Software Implementation

- Testbed： OpenBGPD-8.3-portable, FRRouting-10.2-dev

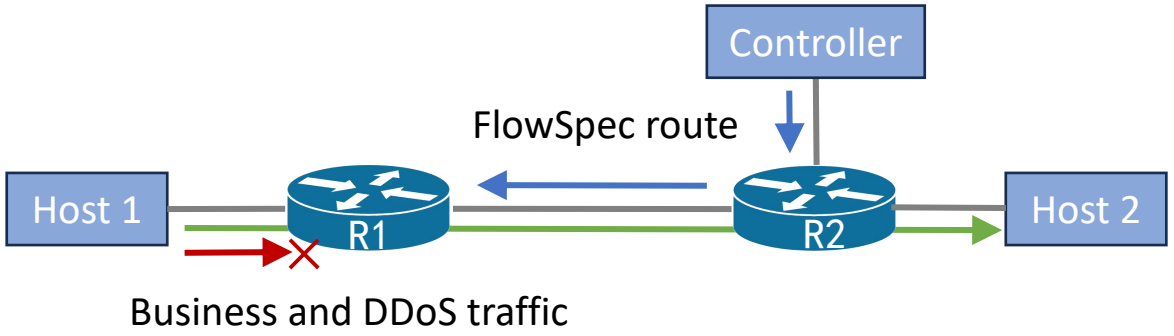

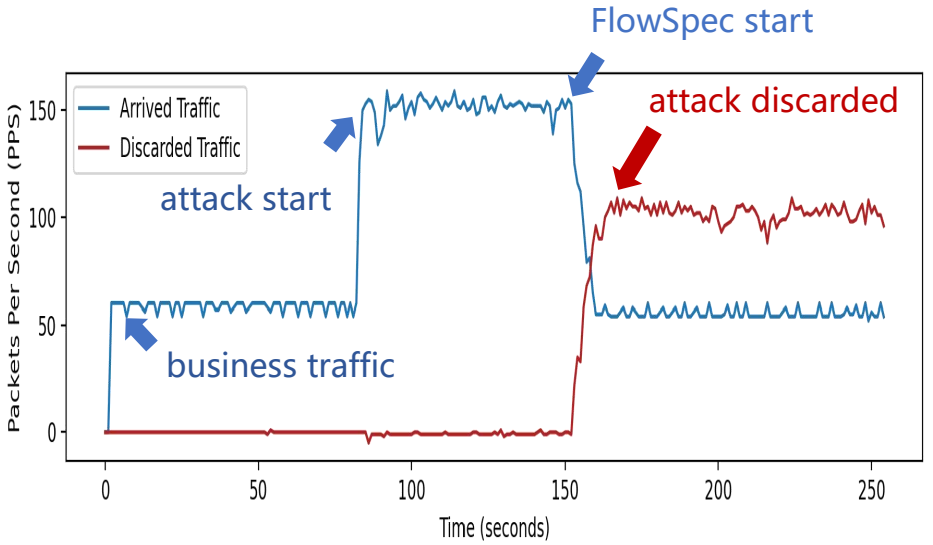

- Defend simple network layer volumetric attack in network device
- Defend 25-55% application layer attack
- Effectively reduce the cost of attack detection and mitigation

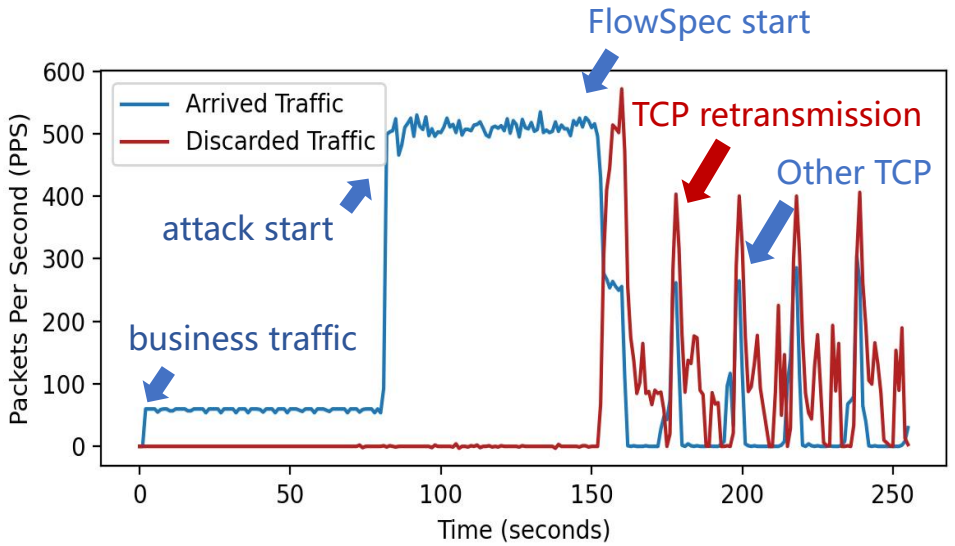

Hardware Deployment

- Testbed： Sino SR3600 switch by ASIC switching chips

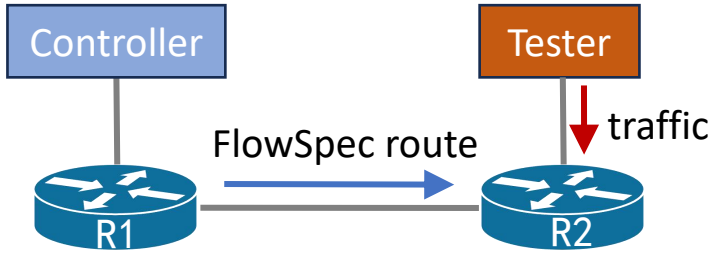


```
SR3600#diagnose
  Enter diagnose mode,please input password:
  %Warning:Executing certain commands may cause abnormal or traffic crash.
           If you need to execute such kinds of commands,Please contact
           the technical engineer,and use it carefully under the guidance
           of the technical engineer.
SR3600(diagnose)#
SR3600(diagnose)#task she
SR3600(diagnose)#task shell srosdriver

***** enter srosdriver shell *****
use "help" for help information

# show_bgp_flow_acl
rule_id    desc                    dir       group_id     entry_id     pri     valid_flag
-------------------------------------------------------------------------------------------
2119       bgp_flowspec            Ingress   2            2119         500000      Valid

 Rule used total:    1

return value = 0x1aa5f98

#
```

- FlowSpec rules are installed as ACL entries
- Payload matching is implemented via UDF and TCAM in hardware
- Rules are enforced at line rate on the ASIC
- Traffic statistics confirm effective packet dropping

```
CTC_CLI(ctc-sdk)# show acl entry-info entry 2119 detail
No.     ENTRY_ID    GROUP_ID    HW  E_PRI       G_PRI   TYPE              INDEX   ACTION    SIZE
-------------------------------------------------------------------------------------------------
1       2119        2           Y   500000      2       Tcam-Flex-100     1280    Discard   320

#ENTRY_ID: 2119          Table:
 --DsProgrammingAclTcam2            :1280
 --DsAcl2Ingress                    :1280
 --DsProgrammingAclTcam3            :1280
 --DsAcl3Ingress                    :1280

Key:
----------------------------------------
UDF             : 00010203040500000000000000000000/FFFFFFFFFFFF00000000000000000000
L3-TYPE         : 2/0xF
IP-PROTOCOL     : 6/0xFF

Action:
----------------------------------------
  DISCARD         : G  Y  R
```

```
25gigaethernet 1/0/44       Down   Full/25G     0.00%    0.00%    0          0
25gigaethernet 1/0/45       Up     Full/10G     10.03%    0.00%    0          0
25gigaethernet 1/0/46       Down   Full/25G     0.00%    0.00%    0          0
25gigaethernet 1/0/47       Up     Full/10G     0.00%    0.00%    0          0
25gigaethernet 1/0/48       Down   Full/25G     0.00%    0.00%    0          0
100gigaethernet 1/1/1       Down   Full/100G    0.00%    0.00%    0          0
100gigaethernet 1/1/2       Down   Full/100G    0.00%    0.00%    0          0
100gigaethernet 1/1/3       Down   Full/100G    0.00%    0.00%    0          0
```

Considerations

- Operational
  - Deploy payload filtering only at controlled locations
  - Enable only on devices that support content parsing and sufficient rule scale
  - Match on the decapsulated inner IP packet to avoid offset ambiguity with MPLS or tunnels

- Security
  - Accept rules only from trusted peers and restrict their propagation scope
  - Combine payload filter with other conditions to reduce false positives
  - Ignore unsupported rules locally and limit update rates to keep BGP stable

- Scalability
  - Limit rules to a small set of high-value entries
  - Use FSv2 rule ordering to limit the scope of traffic inspected
  - Restrict propagation scope to avoid inter-domain scale issues

Next Steps

- Any questions and comments are welcomed
- Optimize fields structure, enhance hardware verification
- Request for WG Adoption

BGP FlowSpec Extension for Feedback Binding

draft-cui-idr-flowspec-feedback-binding-00

Yong Cui, Tsinghua University

Yujia Gao, Zhongguancun Laboratory

Lei Zhang, Zhongguancun Laboratory

May 11th, 2026

Current Draft Status

- 2025/10 first proposed feedback action @IETF 124
- 2026/2 update state machine, and fields design @Interim meeting
- 2026/3 introduce framework and draft scope @IETF 125
- 2026/5 intruduce feedback action @Interim meeting
- 2026/7 Roadmap and next steps to be presented @IETF 126

Problem Statement

- BGP FlowSpec provides one-way control, where successful rule propagation does not imply effective enforcement
- Operators need execution feedback to understand real enforcement behavior and continuously optimize traffic handling strategies

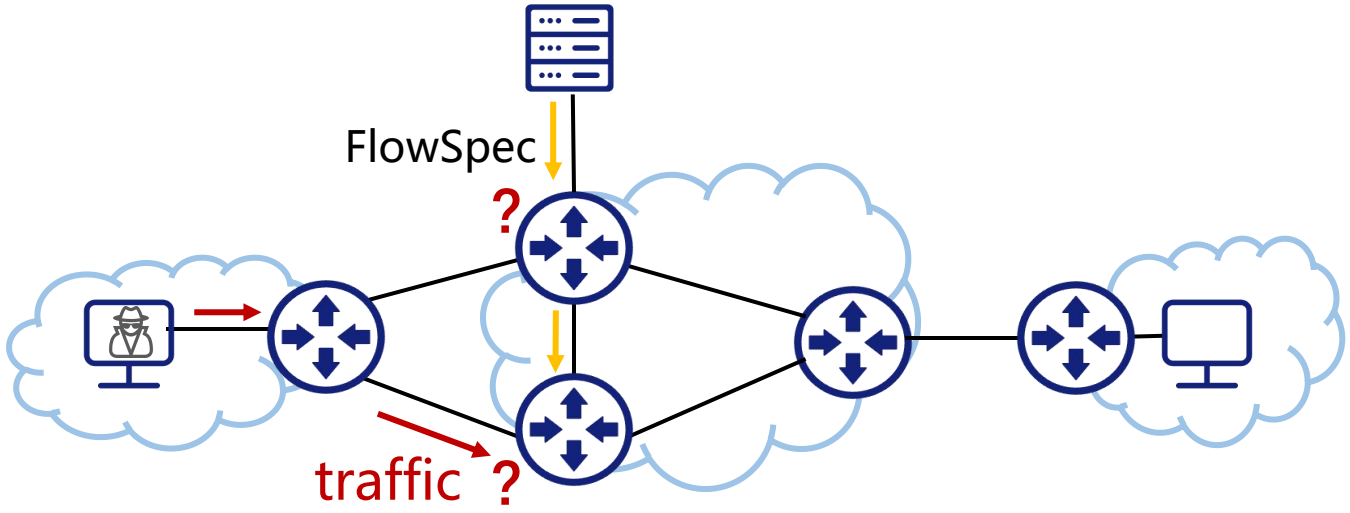

A new FlowSpec action is needed to request execution feedback from receiving devices

FlowSpec Extension of Feedback Action

- Feedback action is encoded in the community container
- Encoding

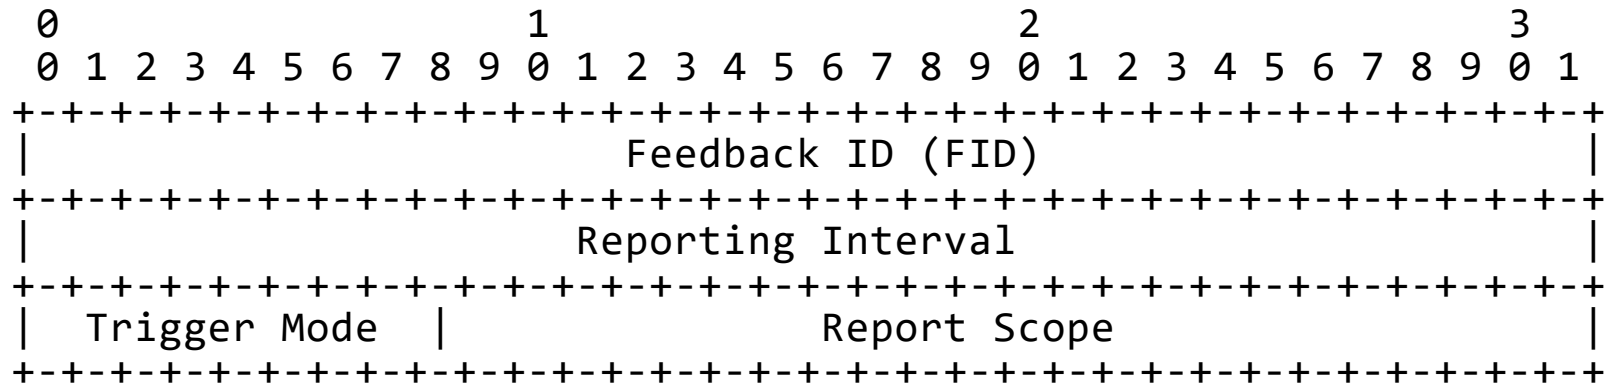

Feedback ID (32 bit): Unique identifier generated by the controller to correlate FlowSpec feedback across rules and devices.

Reporting Interval (32 bit): Periodic feedback reporting interval in seconds; 0 means disables periodic reporting.

Trigger Mode (8 bit): Specifying the feedback trigger mode, 00=periodic enabled, 01=event-only, 10/11 = Reserved

Report Scope (8 bits): field defining the reporting scope, 00=Global, 01=Inter-AS, 10=Intra-AS, 11=Reserved

Next Steps

- Any questions and comments are welcomed
- Clarify the roadmap of the FlowSpec feedback framework
- Optimize fields design and conduct initial validation
- Request for WG Adoption

Thanks!

Contact information: gaoyj@zgclab.edu.cn