Flowspec Capability Bits

draft-haas-flowspec-capability-bits

Jeffrey Haas <jhaas@juniper.net>

Problem Statement (1/2)

- Flowspec serializes packet match rules, typically for firewalls, in BGP NLRI.
- The serialization format in RFC 5575 was «Type, Value» for the variable length items, with the Length portion implied by end-of-list bits. (Compare vs. end-of-stack for MPLS labels.)
- This had the negative consequence that an implementation that didn’t understand a given component (Type) couldn’t safely parse this. Shipping implementations would simply treat unknown types as malformed and would drop the session. This impacts incremental deployment of new Flowspec features.
- This was documented in the update, RFC 8955.

Problem Statement (2/2)

- To address this issue, and perhaps others, work had been started on Flowspec v2. However, that will require a change to Flowspec-encoded NLRI for existing AFI/SAFIs.
  - PCEP had addressed this in its own work (draft-ietf-pce-pcep-flowspec) by leveraging existing Flowspec components, but requiring an explicit Length field in their TLVs.
- Meanwhile, IDR has several pieces of Flowspec work that can’t be incrementally deployed:
  - draft-ietf-idr-flowspec-l2vpn
  - draft-ietf-idr-flowspec-nvo3
  - Others, not yet adopted by Working Group

Proposal

- The issue with current Flowspec is dealing with *unknown* component types. A BGP Speaker can't announce a Flowspec route that a receiver isn't capable of understanding.
- The receiver can use a BGP capability to community which flowspec components it understands.
- This same mechanism can also be used to regulate what a receiver *wants* to receive. It can indicate that it isn't interested in a component type, even if it might understand it.
  - This helps address implementations that have incomplete support for features.

Example encoding

```
Example encoding for Capability Value:

  0                   1
  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |0|1|1|1|1|1|1|1|1|1|1|1|1|1|1|0|0|
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Bit 0 set to 0, bits 1..14 set to 1 showing support for all capabilities for IPv6 Flowspec, bits 15..16 set to 0.

Known Issues

- ”Missing” Flowspec routes can cause unexpected forwarding in a network.
  - The operator must be aware of support for Flowspec features within their network and tune their rules appropriately.
  - This issue already exists when any sort of filtering may be done against Flowspec routes.

Next Steps

- This, or a similar feature, would permit incremental work on Flowspec to continue and get deployed.
- Adopt this draft, or spin something similar to cover the problem space?
- Continue work on Flowspec v2 now that RFC 8955 has shipped.
    - Even in Flowspec v2, this mechanism might be useful, although it’s less critical.

Questions?