# Flowspec Indirection-id Redirect

(draft-vandevelde-idr-flowspec-path-redirect-02)

Gunter Van de Velde, Wim Henderickx, Keyur Patel, Arjun Sreekantiah

# A logical next step for DDoS

- First there was: Indirection to VRF
  - Recursive look-up inside a VRF to find alternate Next-hop destination
  - RFC5575
- Next there was: Indirection to IP
  - Recursive look-up in Routing Table to alternate next-hop
  - draft-ietf-idr-flowspec-redirect-ip
- Now there is: Indirection to Service-plane
  - Recursive lookup to find alternate chain of next-hops
  - Send DDoS traffic over those links/nodes provisioned to transport DDoS
  - draft-vandevelde-idr-flowspec-path-redirect

# Technology Summary

- New BGP Flowspec Action -> redirection using “Indirection-ID”
  - “Indirection-ID” is a new proposed extended community
  - “Indirection-id” is used for a recursive lookup on receiving router
- The flowspec receiving router will use “Indirection-ID” to find out through recursion
  - Tunnel encap information to Next-hop destination
  - Tunnel encap information to Next-Next-hop destination (EPE)
  - Segment Routing or MPLS PCE Binding SID
- A single Flowspec update from a controller results in network wide optimized security, application and traffic steering due to the localized recursion i.e. :
  - Steer to closest IDS or FW or security appliance
  - Steer to engineered path for DDoS mitigation
  - Steer to special regional EPE exit (i.e. for Cloud DDoS handling)
- Steering to Tunnel for DDoS is decoupled from tunnel setup
- Easy to extend extended Community with additional context when use-case prescribes
  - Only ‘B’ Binding tunnel SID defined in draft -0.2
  - Easy to extend to other tunnel mapping contexts

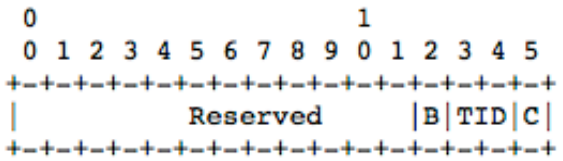

# Use case's covered

- Use-case scenario's:
  - Steer to shortest Path tunnel. Examples could be
    - To regional closest IDS or FW service
    - To best Egress router for the region for flowspec identified traffic
    - To best video rendering device for the region for a particular customer
    - Or simply to a central device in the network
  - Steer to TE-tunnels
    - Steer to RSVP-TE or SR-TE tunnel (to a DDoS mitigation service chain)
    - Steer to Segment Routing binding SID (using a 'bit' in the community local administrator field)
      - i.e. CLI, PCE or BGP based mapping
  - Steer to Next-Next-Hop tunnels
    - Cascaded tunnel tunnel constructs (using "Tunnel ID")
    - Egress Peer Engineering tunnel constructs
    - Engineered path to egress router and exact egress interface

# Comparison Table

| | **draft-vandevelde-idr-flowspec-path-redirect** | **draft-hao-idr-flowspec-redirect-tunnel** | **draft-li-idr-flowspec-redirect-generalized-sid** |
|---|---|---|---|
| Creation Date | **14 September 2015** | 6 October 2015 | 21 March 2016 |
| Airtime at IDR | **IETF94 & IETF95 Interim2015-10-26** | IETF94<br>Interim2015-10-26 | never |
| WG Adoption Call Support(3/25-4/8) | **9 (diverse company support)**<br>(and NONE explicit DO-NOT adopt) | 2<br>(and 2 explicit DO-NOT adopt) | 3<br>(and 1 explicit DO-NOT adopt) |
| Revision | **-02** | -01 | -00 |
| | | | |

# Comparison Table

| | draft-vandevelde-idr-flowspec-path-redirect | draft-hao-idr-flowspec-redirect-tunnel | draft-li-idr-flowspec-redirect-generalized-sid |
|---|---|---|---|
| Next Hop Tunnel support | Yes | Yes | Yes |
| TE tunnel support | Yes | Yes | Yes |
| Nested tunnel support | Yes | No | No |
| Next-Next Hop Tunnel support | Yes | No | No |
| Router Localized tunnel recursion | Yes | No | Yes |
| Tunnel Encap recursion (Flowspec AFI/SAFI coupling with tunnel encap exchange) | Two flavor (IP and non-IP tunnels):<br>**IP:** Decoupled (use-case: SR is not deployed)<br>**None-IP:** SR Binding SID (use-case: SR is deployed)<br>**None-IP:** Decoupled (Use-case: SR is not deployed) | Two flavor (IP and non-IP tunnels)<br>**IP:** draft-ietf-idr-tunnel-encaps<br>**None-IP:** draft-li-idr-mpls-path-programming | **IP:** N/A<br>**None-IP:** Generalized Segment ID ext community contains Segment Routing and tunnel context info (<br>Coupled with draft-li-spring-segment-path-programming) |

# Comparison Table

| | draft-vandevelde-idr-flowspec-path-redirect | draft-hao-idr-flowspec-redirect-tunnel | draft-li-idr-flowspec-redirect-generalized-sid |
|---|---|---|---|
| Flowspec carries tunnel encapsulation | No | Yes (BGP Tunnel Encapsulation Attribute extended to Flowspec AFI/SAFI) (page3, paragraph 1 of draft -01) | No |
| Context extensible | Yes (easy) | No | Yes (harder, assumes using SR or MPP) |
| Use-Case usage | Support for:<br>• shortest path tunnel<br>• TE (rsvp/SR) tunnel<br>• SR binding SID<br>• next-next-hop (EPE, etc..) tunnels | Supported by Section 3.1 in draft-vandevelde-idr-flowspec-path-redirect and hence draft-hao is redundant | Supported by Section 3.2 in draft-vandevelde-idr-flowspec-path-redirect and hence draft-li is redundant |

# WG Adoption Feedback from Ignas Bagdonas

9 April 2016 @ IDR WG Email list:

> Draft-vandevelde can achieve all what draft-hao and draft-li can, and in a more flexible way. Having the ability to decouple redirection tunnel type from redirection action is both practical and extensible - the actual tunnel to be used is a local operational decision for each network element, it is not necessary signalled at the same time and by the same mechanism. Decoupling signalling and redirect parts aligns well to operational practices of using specific tools for specific tasks. Just that BGP could do that does not necesasry mean that it should be used as a best fit. From operational perspective there is no need to have multiple solutions that try to address the narrow problem space in similar yet incompatible ways. There should be one document for redirect, and draft-vandevelde is a good starting base for that.