Strawman for Combining Flow Specification Proposals

Susan Hares

Flow Spec (RFC5575) Review

- NLRI
  - For SAFI 133: IPv4 (AFI=1), IPv6 (AFI=2), L2VPN AFI=25)
  - For SAFI 134: IPv4 (AFI=1), IPv6 (AFI=2), L2VPN (AFI=25)
- Validation
  - Originator of flow spec = originator of best-match unicast route for destination embedded in NLRI,
  - No more specific unicast routes, when compared with Flow destination prefix, that have been received from different neighbor AS
- DoS and L2VPN doesn't fit

  ? – No destination check for SDN/VPN

  ? – does requirement this work for all Filters

  ? ROA's or BGP long-term

BGP Flow Specification is ECA Policy

- **ECA = Event – Match Condition - Action**
  - Flow-specification event = “packet reception”,
  - Condition – match filters in NLRI
  - Action – in Extended communities
- BGP Flow Specifications received
  - (BGP Session ephemeral)

Other ECAs

- I2RS Filter-Based RIB
  - reboot ephemeral
- Policy-Based Routing (filtered routing)
  - configured Flow specification
- ACL
  - Configure flow specification

Why is precedence needed

Precedence needed within BGP Flow Specification

- For filtering
  - For ordering policies - NLRI preference and administrative distance,
    - *Is this BGP normal policies – reasonable or should be able to negotiate (or send in some manner)?*
    - *[Jeff]: Keep deployed FS, Updated Flow specification (address and rule order).*
  - For ordering filters – by Flow Specification type and precedence
- For action
  - No order currently, need to add order

BGP FS Filters types for RFC/WG documents

- RFC 5575 types/v6-draft
  1. Destination prefix
  2. Source prefix
  3. IPv4 protocol / IPv6 Next header
  4. Port (source or destination)
  5. Source port
  6. Destination port
  7. ICMP Type
  8. ICMP Code
  9. TCP Flags
  10. Packet length
  11. Traffic Class
  12. IPv4 Fragment
  13. IPv6 Flow ID
- L2VPN types
  14. Ethernet type
  15. Source MAC
  16. Destination MAC
  17. DSAP in LLC
  18. SSAP in LLC
  19. Control fields in LLC
  20. SNAP
  21. VLAN ID
  22. VLAN COS
  23. Inner VLAN ID
  24. Inner VLAN COS

BGP FS Proposed Filter types

- MF-1 NV03 Delimiter
  - Inner/outer header info
- MF-2 Virtual Network ID (VNID)
- MV-3 Flow ID (NVGRE Flow ID)

Missing

- MF-4 MPLS LSP label or label stack
- MF-5 Interface Grouping
- MF-6 Time matches

Are there others?

- Other types?
- Should we set a few types, and then create an Extended BGP Flow specifications
  - In another NLRI,
  - Or another BGP Attribute
  - (draft-li-flowspec-rpd)

Precedence Rules

Precedence logic for BGP Flow Specifications
(RFC5575, draft-idr-bgp-flowspec-l2vpn)

```
flow-rule-cmp (a,b)
{
  comp1 = next_component(a);
  comp2 = next_component(b);
  while (comp1 || comp2) {
   // component_type returns infinity on end of list
   if (component_type(comp1) &lt; component_type(comp2)) {
    return A_HAS_PRECEDENCE;
    }

   if (component_type(comp1) &gt; component_type(comp2)) {
    return B_HAS_PRECEDENCE;
   }
```

Precedence Rules (2)

```
// IP values)
  if (component_type(comp1) == IP_DESTINATION || IP_SOURCE) {
    common = MIN(prefix_length(comp1),prefix_length(comp2));
                  cmp = prefix_compare (comp1,comp2,common);
                  // not equal, lowest value has precedence
                  // equal, longest match has precedence;
  } else if (component_type (comp1) == MAC_DESTINATION ||
          MAC_SOURCE) {
                                common = MIN(MAC_address_length(comp1),
                                        MAC_address_length(comp2));
                                cmp = MAC_Address_compare(comp1,comp2,common);
                                //not equal, lowest value has precedence
                                //equal, longest match has precedence
                  } else {
     common = MIN(component_length(comp1),
                                        component_length(comp2));
                   cmp = memcmp(data(comp1), data(comp2), common);
                                //not equal, lowest value has precedence
                                //equal, longest string has precedence
   }
  }
}
```

Flow Specification Actions

Approved Actions (RFC 5575 and RFC 7674)

- Traffic rate in bytes (0x8006)
- Traffic Action (0x8007) with S(sample) T (terminal) flags
- Redirect to IP VPN via Route Target
  - RD 2 octet AS, 4 byte value (0x8008)
  - RD 4 octet IP, 2 byte value (0x8108),
  - RD 4 octet AS, 2 byte value (0x8208)

Proposed Actions

- (FA1) Traffic Rate in packets
- (FA2) Traffic Action with "R" for refer to more policy in BGP Attribute
- (FA3) Redirect to Tunnel
- (FA4) VLAN Action
- (FA5) TPID action
- (FA6) MPLS label action (push, pop, swap)
- (FA7) interfaces set
- (FA8) change validation to ROA or bgpsec-protocol

Default Precedence Proposals for BGP Flow Specification

- Filters – AND
  - Change filter order [MF-0/MF-xx]
  - IP Protocol (1-13)
  - NVO3 matches (MF1-MF3)
  - L2VPN matches (14-24)
  - MPLS matches (MF-4)
  - Interfaces matches (MF-5)
  - Time matches (MF-6)

Action

1. Change order of actions (FA-0)
2. Traffic rate in bytes
3. Traffic rate in packets (FA-1)
4. Traffic Action (RFC5575)
5. Extended Traffic Action (FA-2)
6. Redirect to IP VPN
7. Redirect to tunnel (FA-3)
8. VLAN action (FM-4)
9. TPID action (FM-5)
10. Label Action (FM-6)
11. Interface Set (FM-7)
12. Use Alternate BGP Validation

Possible Conflicts

	Possible conflicts										
Action	Traffic rate Bytes	Traffic Rate Pkts	Traffic Action	Ext. Traffic Action	Redirect To IP VPN	Redirect to IP Tunnel	VLAN	TPID	Label	Intf Set	BGP valid
Traffic Rate Bytes		X									
Traffic rate Pkts	X										
Traffic action				X							
Ext. Traffic action			X								

Possible Conflicts

	Possible conflicts										
Action	Traffic rate Bytes	Traffic Rate Pkts	Traffic Action	Ext. Traffic Action	Redirect To IP VPN	Redirect to IP Tunnel	VLAN	TPID	Label	Intf Set	BGP valid
Redirect IP VPN						X	X	X	X	X	
Redirect Tunnel					X		X	X	X	X	
VLAN					X	X		X	X	X	
TPID					X	X	X		X	X	
Label					X	X	X	X		X	
Intf. Set					X	X	X	X	X		

BGP Flowspec vs. I2RS Filters

component	BGP Flow Spec ECA	I2RS FB-RIB Packet-ECA
Policy	flowspec-policy	rule-group
+-name	+-policy-name	+-group-name
+-vrf	+-vrnf-name	
+-AFI	+-address-family	
+-rule	+-flowspec-rule [rule-name]	+rule[rule-name]
+-name	+-rule-name	+-rule-name
+order		**+-rule-order**
+installer		+-installer
+-match-filter	+--flowspec-component	+-eca-match-filter
+-type	+-component-type	+-match type
+-intf	+-Intf Group	+intf

Bgp Flow Spec vs I2RS Filters

component	BGP Flow Spec ECA	I2RS FB-RIB Packet-ECA
More filters		
+-intf	+-Intf Group	+intf
+-L1		+l1 header
+-L2		+L2 header
+-MAC		+-MAC
+-L2VPN	+-l2VPN	+-L2VPN
+-MPLS	+-MPLS	+-MPLS
+-Nvo3	+-NV03	+-NVO3
+-L3	+-L3	+-L3 header
+dst IP	+dst IP	+dst-ip
+src IP	+src IP	+src-ip
+proto	+proto	+proto
+dscp	+dscp	+dscp
+ICMP	+icmp type	+icmp type
	+icmp value	+icmp code
+v6flow	+v6-flow-id	+v6flow
+-L4	+L4	+-L4-header
+S-port	+src port	+src-port
+d-port	+dst port	+dst-port

BGP FS vs. I2RS Filters Yang

component	BGP Flow Spec ECA	I2RS FB-RIB Packet-ECA
… operational state		
+-ro flow- state ...	+ro-flowspec- state ...	+ro fb-rib opstate (need to add) ...
Operational Statistical State		
+-ro flow- stats	+ro-flowspec- stats	)(need to add)

Discussion

- Should we have a successor to Flow-spec SAFI?
  - Action Criteria: IP Redirect (do-able) with 2 feature; Combination become with Actions is tricky;
    - Choice: combination
    - Precedence: better to specify, but will need to consider actions in combination
    - Redirect actions – interact with each; Modify actions interaction;
    - Traffic filters may
  - Match filters – as AND probably

Discussion (2)

- Flow Specification
  - Combination or separate Flow Spec
  - Rule ordering is reason for Flow-Spec 2,
    - Non-firewall, no SDN –may work
    - Firewall, SDN will not work without the ordering
  - Combination of the two flow-specification
    - If keep 2 SAFIs – two Flow-Specs into the future.
    - Ideal, v2 would have package with it – Date to deprecate V1 – real world doesn't probably won't allow it,
  - Agree with Jeff on backward compatibility
    - [wes] No way to tell which enhancement supported with out pre-knowledge,
    - [Jeff]: We do not have way to discover capabilities
    - [Robert]: We have this problem with the
    - [Jeff]: Redirect IP – possible that flow-specification – action (what does the implementation do with it).
  - Inter-domain flow-specification – not common
    - Service portals rather than inter-AS Flow specification
    - Redirect IP – within a single Provider – within a specific Provider

Discussion (3)

- Centralized mode –
  - Some flow specifications are only centralized controller and not distributed (Lucy Yong)
  - Some have two controllers (DDoS) and another (flow-filters)
    - Need to have precedence of the rules and then fall through (Jeff)
    - SDN (rule), and then the flow-specification rule
    - This requires a flow-specification v2 (jeff) because the existing things do not allow the flow-specification
    - Some the actions may only be appropriate to the list
  - Filter-based RIB
    - Precedence, fall-through – rule chains make sense
    - Take I2RS Filter-Based RIB

Discussion (4)

- Implementation of I2NSF
  - Controller tells the order of the rules
  - Can IDR provide this as well.
  - [Jeff]: More specific hosts, flow specification (longest prefix match will work)
  - [Linda]: Most specific

Discussion (5)

- John Schiel – flow spec rules that have precedence and ordering in flow specification rules.