

# Dissemination of Flow Specification Rules

**draft-ietf-idr-rfc5575bis-09**

## Status Update

**Susan Hares**

Huawei

**Christoph Loibl (speaker)**

next layer Communications

**Robert Raszuk**

Bloomberg LP

**Danny McPherson**

Verisign

**Martin Bacher**

T-Mobile Austria

# Motivation

- Lack of interoperability
  - Unclear sections in the specification (hard to implement consistently)
- Unspecified behaviour on interfering flow-actions
- Preserve compatibility - if possible

# Status

- WG LC shortly after IETF 99
- Minor issues found (-07, -08) on Jie's shepherd review
- Minor issues found (-09) on Jeffrey's review
- Jeffrey also pointed out issues that we think should be discussed

# #1 BGP Opaque NLRI(or not?) (in RFC 5575 and in -bis)

BGP itself treats the NLRI as an **opaque key to an entry in its databases**. Entries that are placed in the Loc-RIB are then associated with a given set of semantics, which is application dependent.

< ... >

Standard BGP policy mechanisms, such as **UPDATE filtering** by NLRI prefix and community matching, **SHOULD apply** to the newly defined NLRI-type. Network operators can also control propagation of such routing updates by enabling or disabling the exchange of a particular (AFI, SAFI) pair on a given BGP peering session.

RFC 5575 - Section 3

# #1 BGP Opaque NLRI (in RFC 5575 and -bis)

- Opaque (keep as is):
  - Propagate garbage
  - Propagate unknown FS components (simple extension of the FS NLRI)
- Non-Opaque (remove opaque property from draft):
  - Known BGP implementations “Non-Opaque”
  - More consistent within the protocol (validation, BGP update filters)
  - Changes to the draft primarily to allow future FS-NLRI extensions
    - Treat as withdraw + always “true” match condition for future NLRI extensions

# #2 Very restrictive -bis Section 7.6.

## Rules on Traffic Action Interference

- Traffic actions encoded as extended communities
- Current draft: **resolve conflicts -> treat as withdraw**
  - Example: 1+ rate-limits, 1+ redirect-rt
  - esp. multiple redirect-rt actions may make sense when different routing information is available at different nodes in the network
- Suggested less restrictive: Having a **predictable behaviour** can also be achieved by some **“sorting”**
  - ie “lowest” redirect community possible gets applied (plus some default if non of the redirect communities can be applied.
  - ie “lowest” rate-limit gets applied

# Implementation reports

- Most implementations of RFC 5575 comply to a very high degree also with -bis draft (even more if the opaque key property gets removed)

# Questions to WG

- Remove Opaque property from the draft and resolving extensibility issues by “treat as withdraw”++?
- Make interfering actions less restrictive and achieve predictability by sorting them?