

# Carry congestion status in BGP extended community

Zhenqiang Li (editor)

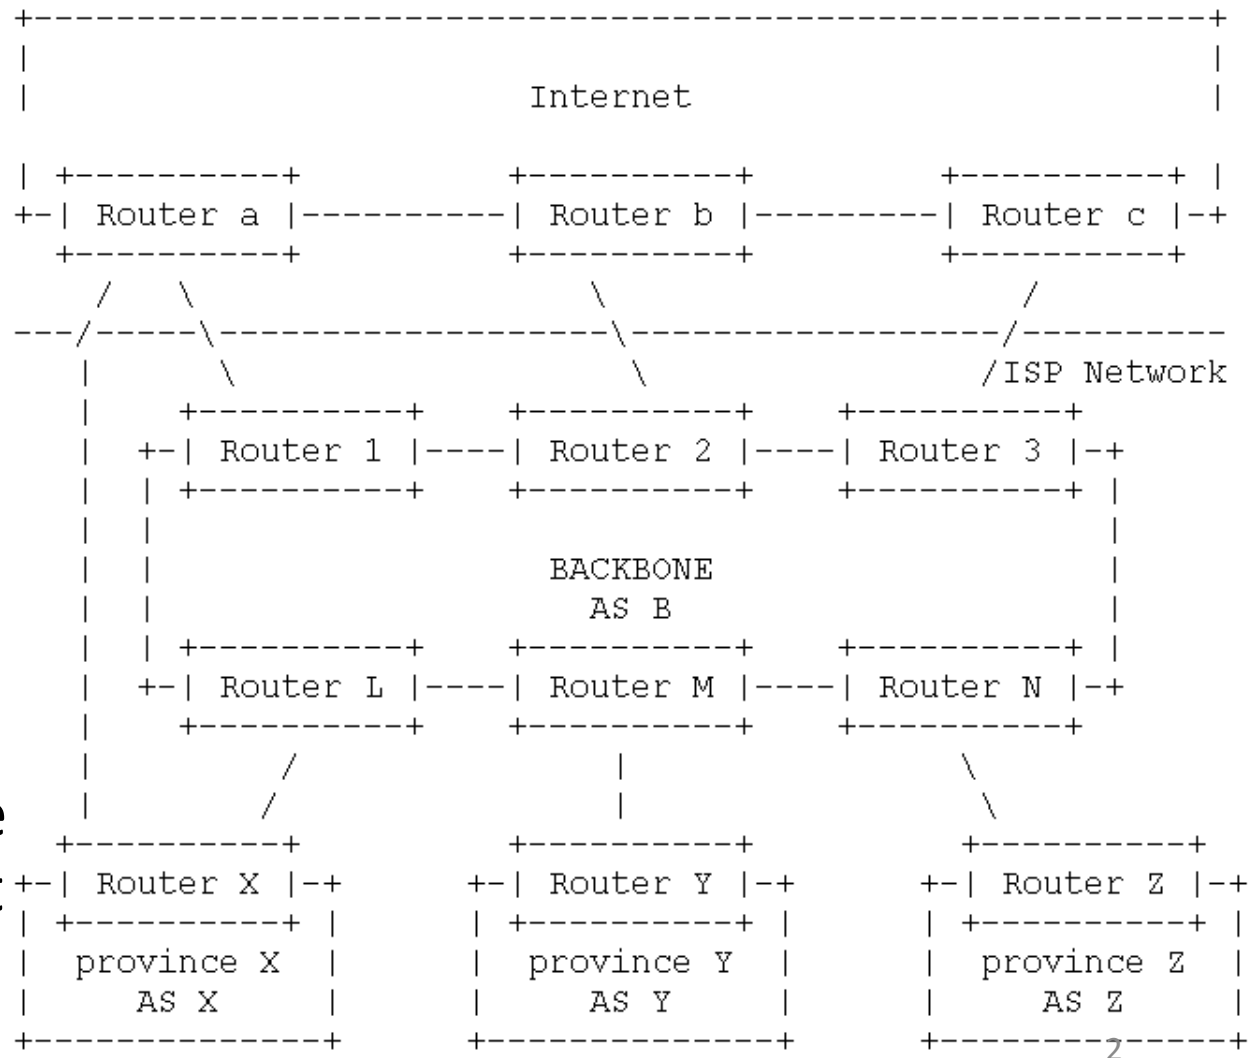
China Mobile

Jie Dong

Huawei Technologies

# Scenario to be addressed


- The congestion status of the exit links in one AS can be used by the BGP receivers to steer the traffic going out the AS through route policy.
- This document introduces a new extended community to delivery the congestion status of the exit link to other BGP speakers.
- Congestion status extended community is good not only to the ASBRs in other AS, but also to the BGP peers within one AS.



# Congestion Status Extended Community


- Since Congestion status extended community can be used by BGP speakers in other AS, it MUST be transitive, i.e. the T bit in the first octet of extended community MUST be zero.
- Congestion status extended community has two encoding formats, one is for two-octet AS, the other is for four-octet AS.

# Congestion Status Extended Community for Two-Octet AS



- It is a sub-type allocated from Transitive Two-Octet AS-Specific Extended Community Sub-Types defined in section 5.2.2 of [RFC7153].
- The "Type" field MUST be 0x00 to indicate this is a Transitive Two-Octet AS-Specific Extended Community.
- The "Sub-Type" field is used to indicate this is a Congestion Status Extended Community. Its value is to be assigned by IANA .
- The "Sender AS Number" field stores the AS number of the BGP speaker who generates this community.
- The "Bandwidth" field is 1 octet. Its value is the bandwidth of the exit link in unit of gbps.
- The "Utilization" field is 1 octet. Its value is the utilization of the exit link in unit of percent.

# Congestion Status Extended Community for Four-Octet AS



- It is a sub-type allocated from Transitive Four-Octet AS-Specific Extended Community Sub-Types defined in section 5.2.4 of [RFC7153].
- The "Type" field MUST be 0x02 to indicate this is a Transitive Four-Octet AS-Specific Extended Community.
- The "Sub-Type" field is used to indicate this is a Congestion Status Extended Community. Its value is to be assigned by IANA .
- The "Sender AS Number" field stores the AS number of the BGP speaker who generates this community.
- The "Bandwidth" field is 1 octet. Its value is the bandwidth of the exit link in unit of gbps.
- The "Utilization" field is 1 octet. Its value is the utilization of the exit link in unit of percent.

# Security Considerations

- Malicious router may use the congestion status extended community to interfere the traffic steering decision of the BGP receiver.
- BGP peers SHOULD use MD5 for authentication [RFC4360]. BGP receiver SHOULD only accept the congestion status extended community delivered from BGP peers with MD5 authentication.

# IANA Requirements

- For Congestion Status Extended Community for Two-Octet AS, one sub-type is solicited to be assigned from Transitive Two-Octet AS-Specific Extended Community Sub-Types registry.
  - 0x06 is suggested.
- For Congestion Status Extended Community for Four-Octet AS, one sub-type is solicited to be assigned from Transitive Four-Octet AS-Specific Extended Community Sub-Types registry.
  - 0x06 is suggested.

- Thanks
- Comments to
  - [li\\_zhenqiang@hotmail.com](mailto:li_zhenqiang@hotmail.com)
  - [idr@ietf.org](mailto:idr@ietf.org)