IDR Interim
BGP Autoconfiguration

June 21, 2021

Agenda

- Scope of discussion
- State carried in the protocol
- Transport and BGP protocol considerations
- What protocols do we create from this state?
- Security considerations for the auto discovery protocol

Scope of discussion

- Reminder: Initial focus is data center
- While most of the state is similar for multi-hop BGP (internal or external), we may require additional auto-discovery state. For example:
  - TTL
  - Path MTU Discovery settings

State needed by auto-discovery

- BGP Session Transport State:
  - IP addresses
  - Transport security parameters
  - GTSM [RFC5082] configuration, if any
  - BFD [RFC5880] configuration, if any BGP
- Session Protocol State:
  - AS Numbers
  - BGP Identifier
  - Supported AFI/SAFIs
  - Device Role (future extension?)

How to do extensions?

State needed by auto-discovery (2)

Session protocol state is capable of being "Discovered at BGP Open" if you connect to the BGP peer.

- This avoids potentially conflicting state.
- It means the only way for a client to figure it out is to connect.
- Impact point is how often peers try to connect (perhaps repeatedly) to devices announcing auto-discovery.
  - This can be mitigated by putting information that state has changed.
- Router servicing incoming session that reaches Established spends resources for operating BGP that may be immediately discarded if the discovering device decides that the session is unacceptable.

- Session Protocol State:
  - AS Numbers
  - BGP Identifier
  - Supported AFI/SAFIs
  - Device Role

Transport and BGP protocol considerations

- In order for BGP to be able to succeed for auto configuration, the BGP TCP session must be able to come up:
  - IP Endpoints must be known.
  - GTSM must be consistently applied, if used.
  - Authentication or transport security needs to be consistent
  - Once BGP comes up, if BFD procedures are inconsistent, session won't survive. This can be obviated by draft-ietf-idr-bfd-strict.
- BGP's state machine can handle starting the connection from auto-discovery as part of a "Manual Start". However, on failure, retry timers may be inappropriate for auto-discovery environments.
  - Aggressive timers may be problematic, especially at scale.
  - "At scale" may not apply to the DC case.

What protocols do we create from this state?

- The primary consideration for the design team was “for data centers”.
  - The same mechanism is likely applicable for more general cases.
- Layer 2
  - It doesn’t route
  - It’s on the same link
  - Security and privacy considerations are possibly constrained
- Layer 3
  - Likely requires IP multicast

Security considerations for the auto-discovery protocol

- Auto-discovery doesn’t bypass security mechanisms on BGP sessions.
- It however can potentially trigger aggressive BGP connection attempts on a BGP implementation.
  - At scale, this is a denial of service issue.
- Security ADs will likely require protocol to carry some minimal authentication/integrity information.
  - One authentication profile may be “NULL”.