

Progress of BGP Autoconf Design Team

Jie Dong

(on behalf of the design team members)

Purpose of Design Team

- The team is initially chartered to consider requirements and review the solution space (notably including existing proposals).
 - What we would like from the design team is **proposed requirements** and **solution space outline** to guide our discussion as a WG in producing a solution (which will likely include one or more interim meetings).
 - The design team isn't forbidden from suggesting a solution, but it's not one of our specific requests to them.

Design Team Members

- Acee Lindem
- Jeff Haas
- Jeff Tantsura
- Jie Dong (Lead)
- Mankamana Mishra
- Randy Bush
- Robert Raszuk
- Warren Kumari
- Xiaohu Xu

How We Work

- Discussion mostly happens on the design team mail list:

<https://mailarchive.ietf.org/arch/browse/bgp-autoconf/>

- Held one conference call to speed up the discussion
 - Minutes: <https://etherpad.ietf.org:9009/p/bgp-autoconf-feb-25>
 - Some difficulty in finding appropriate time slot for all members

The Scope

- The design team has agreed to work on the DC case first
 - More precisely, when BGP is used as the underlay routing protocol in data center
- Will also keep an eye on the difference in other cases
 - WAN
 - IXP
 - ...

Requirements Collected (1): Common Ones

- Support IPv4 and IPv6 address family
- Support to use either interface or loopback address for BGP session
- Support to discover the peering IP address
- Support to discover the peering ASN
- Support authentication of control message
- Enable Layer 3 link liveness detection, such as BFD

Requirements Collected (2): Under Discussion

- The capability of communicating arbitrary attributes to peers according to operators' need
 - The information should be only sent to peers, and not propagate further
- While it is considered useful, there is ongoing discussion about:
 - Whether it should be part of BGP autoconf, or it can be done in BGP itself (i.e. after the session is established)
 - Whether it should be a generic capability for operator's customization, or some guidance or structure needs be specified as part of the design?

List of Other Suggested Features (1)

- These features are NOT adopted yet but may be considered
 - Discover mutually supported encapsulation
 - Provide Layer 2 keep-alive messages for session continuity
 - Discover the role of the connected nodes
 - Automatic setup of reachability to peer's loopback over one or more connected links
 - Provide resolution for the BGP next-hop address (i.e. the loopback address) for the BGP routes exchanged over these sessions between the loopback addresses.
 - Enable exchange of IP addresses and link attributes between the directly connected BGP routers. should be extensible to include other information in future.

List of Other Suggested Features (2)

- These features are NOT adopted yet but may be considered
 - Discover neighbor's BGP ID for consistency check or avoid connection collision
 - Discovery parameters relating to the BGP peer session (e.g., the local address)
 - Mechanism should be limited to link scope for security and use link-local addressing only
 - Support optional validation of parameters to detect misconfiguration (e.g. link address subnet mismatch, peering between incorrect AS, etc.) in an extensible manner

Design Principles: Needs Further Discussion

- Independent from BGP session establishment
- Not affect or change BGP session establishment and routing exchange, other than the interactions for triggering the setup/removal of peer session based on discovery mechanism
- Generic for any link-layer protocol
- Make use of a currently implemented and deployed DC switch protocol to reduce the cost and complexity
- Make use of existing BGP protocol for automating the BGP session bring-up
- Widely applicable to a range of routing and similar protocols which need layer 3 discovery and characterization
- Length of the message size supported

Existing Proposals

- draft-acee-idr-ldp-peer-discovery-06
- draft-xu-idr-neighbor-autodiscovery-12
- draft-ymbk-lsvr-l3dl-ulpc-02
- draft-raszuk-idr-bgp-auto-session-setup-01
- draft-raszuk-idr-bgp-auto-discovery-06
 - Merged draft-raszuk-idr-ibgp-auto-mesh-00 and draft-wkumari-idr-socialite-02
 - Applicable to BGP auto discovery in WAN & IXP
 - Not included in the analysis on next page

A Brief Analysis to Existing Proposals

Draft Name	Design Principle	Basic Functions	Extensibility	Concerns
draft-acee-lldp-peer-discovery-06	Extensions to LLDP (layer 2)	Discover peering address, ASN, BGP ID, Group ID , capabilities, key-chain, local address; Support BGP session on direct / loopback*	Easily extensible with additional TLVs	LLDP message length limitation; Progress of LLDPv2
draft-ymbk-lsvr-l3dl-ulpc-02	Extensions to L3DL (layer 2)	Discover peering address, ASN, authentication data; Support BGP session on direct/ loopback*	Easily extensible with additional TLVs	Rely on the support of L3DL, session based
draft-xu-neighbor-autodiscovery-12	New BGP message based on UDP	Discover peering address, accepted ASN list, Local prefix, Link attribute , authentication; Support BGP session on direct/ loopback IF, support ECMP;	Easily extensible with additional TLVs	Change to BGP, additional FSM
draft-raszuk-idr-bgp-auto-session-setup-01	Reuse BGP OPEN message with new UDP port	Discover peering IP address, ASN, BGP ID; Support BGP session on direct/ loopback	Constrained with BGP OPEN optional parameters*	Solution not quite complete

* Not fully specified in current draft

Next Steps

- Confirm the minimal set of common requirements
- Reach consensus on the design principles
- Probably put them into a requirement document
- Hand the solution discussion to the WG

Thank You