

BGP Peer Discovery Using IEEE LLDP IETF 99, Prague

Acee Lindem, Cisco
Keyur Patel, Arrcus
Shawn Zandi, Linkedin
Jeff Haas, Juniper
Xiaohu Xu, Huawei



BGP Peer Discovery Requirement



- Support BGP peer discovery for on switch/router ports
- Minimal configuration differences between switch configurations
 - Initial Requirement for Data Center Topologies
- Support for peering on loopback addresses
- Support for discovery of authentication methods and mechanisms

BGP LLDP Peer Discovery



- Use LLDP (IEEE 802.1AB) for BGP Peer Discovery
- IANA already has an IEEE OUI and LLDP allows advertisement of Organizationally Specific TLVs in LLDP PDUs (Protocol Data Units).
- Add IANA Registry for IETF LLDP OS-TLVs
- First IETF LLDP OS-TLV will be BGP Config
- LLDP BGP Config OS-TLV will support Sub-TLVs
 - Peering Address (for one or more <AFI, SAFI> tuples)
 - Local AS
 - BGP Identifier
 - BGP Group ID
 - BGP Session Capabilities
 - Key-Chain Name

Changes Since Chicago

IDR (1/2)



- BGP's LLDP usage is purely as a discovery protocol
 - No longer need to advertise interface addresses in LLDP (only peering addresses)
 - No longer need to add and delete routes to 2-hop peering addresses
 - 2-hop peering is still supported but reachability is out of scope
- Added peering address wildcard <AFI, SAFI> tuple for simplification in deployments with uniform address family usage.
 - Caveat with respect to capabilities negotiation for non-uniform deployments

Changes Since Chicago

IDR (2/2)



- Add ability to advertise an optional second local AS
 - Can be used for AS transition scenarios
- Addition of optional BGP Identifier Sub-TLV
 - Can be used for validation and to reduce possibility of connection collisions
- Addition of BGP Group-ID
 - Abstraction used to classify supported session type
 - For example, in the data center it could be used to classify ToR or Spine BGP sessions
- Addition of optional key-chain name Sub-TLV
 - Can be used to identify keys for MD5 or TCP-AO authentication

Next Steps



- Gauge Interest over proprietary and more limited mechanisms (e.g., ICMPv6 Router Advertisements).
- Request WG Adoption

LLDP, IEEE 802.1AB in a Nut Shell



- Unidirectional Layer 2 discovery protocol with minimal state.
- 3 Mandatory TLVs Chassis ID, Port ID, and TTL.
 - Chassis ID and Port ID tuple defined MAC SAP with which to associate advertised information
 - TTL defines the life of the advertised information.
- Optional TLVs including TLVs defined by other organizations (Organizationally Specific TLVs or OS-TLVs).
 - All IETF needs to advertise is an IANA registry since IANA has its own OUI
- Transmission initiated based on timer expiration or LLDP PDU TLV value change.
 - Very simple but limits amount of information that can be advertised to single LLDP PDU
 - LLDP PDU information replaces the previous information