BGP Peer Discovery Using IEEE LLDP IETF 98, Chicago

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

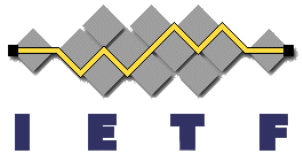

BGP Peer Discovery Requirement

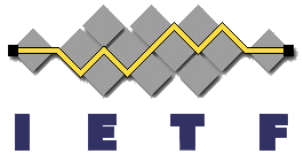

- Support BGP peer discovery for on switch/router ports
- Minimal configuration differences between switch configurations
  - Initial Requirement for Data Center Topologies
- Support for peering on loopback addresses
- Support for discovery of other authentication methods and mechanisms

Why Link Layer Discovery Protocol (IEEE 802.1AB)?

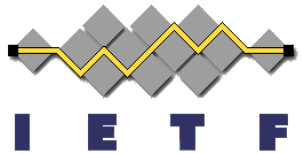

- Discovery protocol already implemented on many, if not most, commercial routers and switches.
- Non-proprietary IEEE standard which already supports advertisement of information from other organizations (e.g., IETF). IANA already has their own OUI.
- Leveraging this widely deployed protocol for IETF discovery is viewed as generally useful by, at least, the authors
- Unlike IGPs, BGP has no native discovery program.
- IPv6 Neighbor Discovery solution is limited to peering with any neighbor on switch port.
  - Must peer on link-local address and use RFC 5549 for IPv4 prefixes.

LLDP, IEEE 802.1AB in a Nut Shell

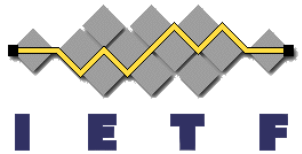

- Unidirectional Layer 2 discovery protocol with minimal state.
- 3 Mandatory TLVs Chassis ID, Port ID, and TTL.
  - Chassis ID and Port ID tuple defined MAC SAP with which to associate advertised information
  - TTL defines the life of the advertised information.
- Optional TLVs including TLVs defined 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

New LLDP Organizationally Specific TLVs - (OS-TLVs)

- Interface Address OS-TLV – IPv4 or IPv6 address on the layer 2 interface.
  - More than one may be advertised.
- BGP Peering Address OS-TLV – Supports multiple Sub-TLVs relating to BGP Peering
  - Local-AS Sub-TLV – 32-bit local AS for BGP peering
  - Peering Address Sub-TLV - Addresses and corresponding AFI/SAFI Tuples
    - Supports different peering addresses for different Address Families (characteristic of previous auto-discovery approaches)
    - All zeros AFI/SAFI is wildcard indicating any MP-BGP negotiated Address Families
  - Capabilities Sub-TLV – Currently MD5 authentication, TCP-AO authentication, and GTSM

Protocol Operations – Work in Progress

- LLDP Peering Addresses used for BGP session establishment
  - BGP sessions are NOT torn down based on changes
- Routes will be dynamically added for 2-hop session loopbacks
  - Can be ECMP for loopback if same peering address received on more than one port
  - Withdrawal of interface address or change in peering address will result in route change and possible unreachability for peering address
    - Don’t see alternative for this feature
- GTSM TTL will be determined dynamically based on 1 or 2 hop session (thanks John Heasley)
- Discussions as to whether authentication come with some SA (Security Association) ID

Next Steps

- Discussion and Gauge Interest
- Progress based on interest