Advertising TE Information in BGP

draft-gredler-bgp-te-00

Hannes Gredler <hannes@juniper.net>
Jan Medved <jmedved@juniper.net>

Motivation (1)

- Look across the “fence”
  - “Fence” being IGP area/level or AS boundary
- Gain visibility for application(s) which need ***complete*** topology data

Use case Alto Servers: multi-area IGP topology

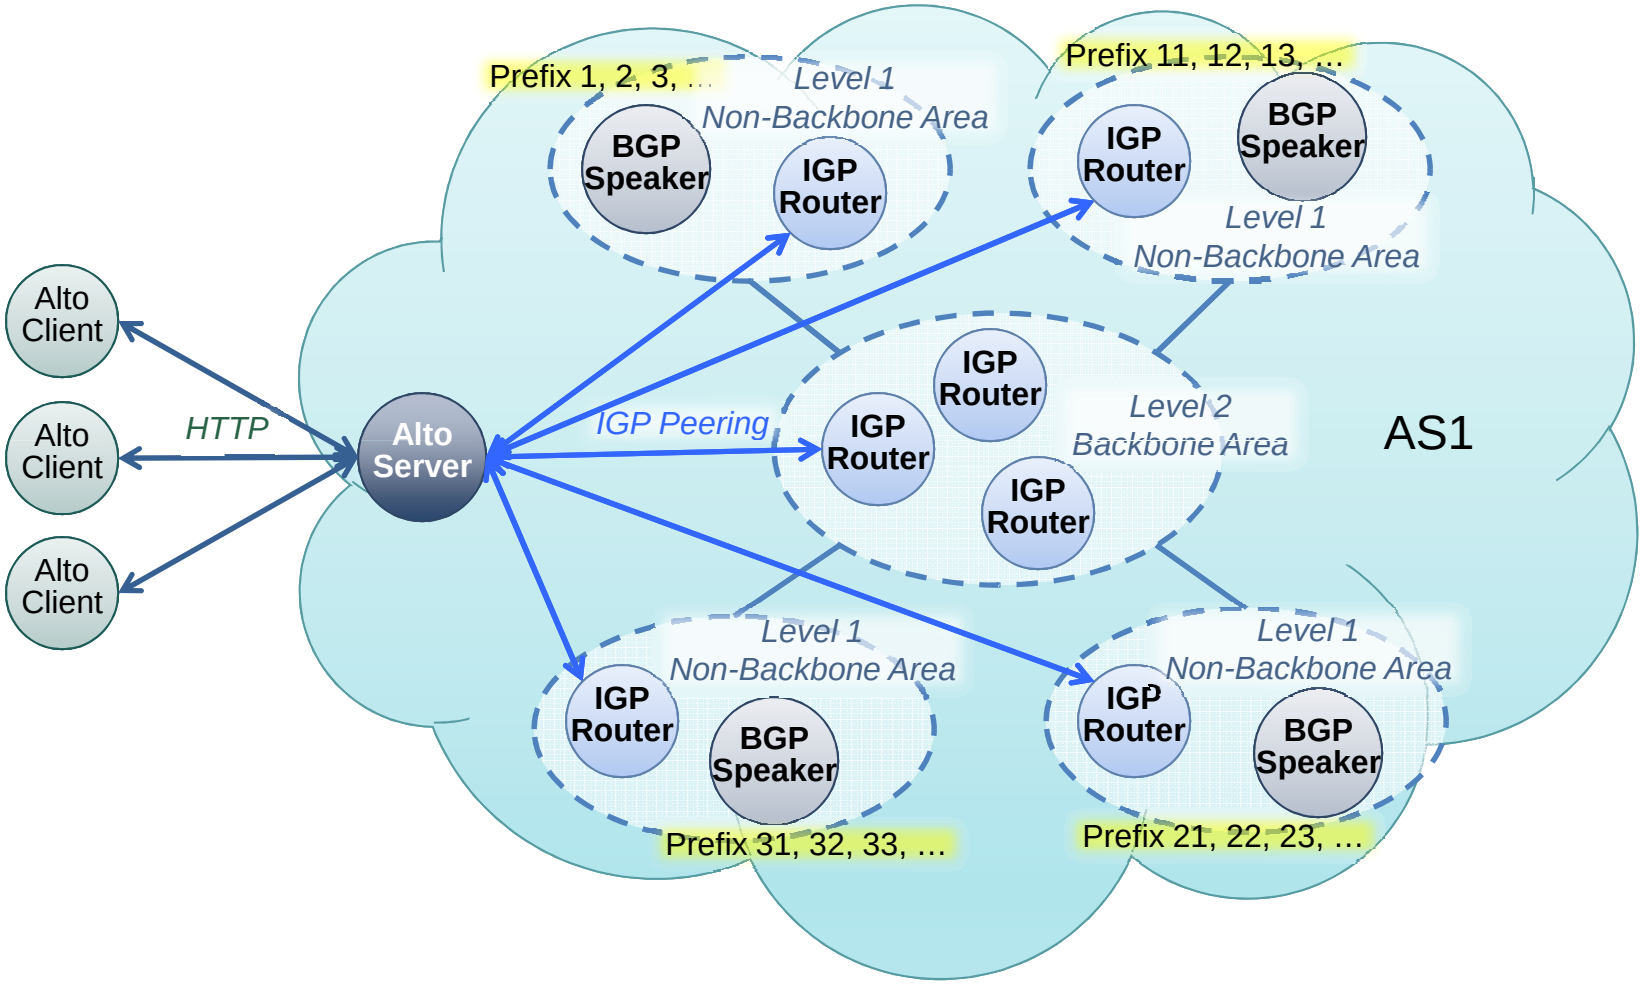

- ALTO server needs to know all areas topology
- Manually crafting of “IGP peering” topology is tedious and error prone

Motivation (2)

- Knowing remote topology & TE data from a remote area useful for:
  - MPLS TE: visibility into remote areas/ASes enables diverse paths through the entire network
  - PCE TED synchronization protocol
  - Network API to communicate topology to outside servers (ALTO Server, Path Computation Server)
- Scope:
  - Attributes/metrics of links/paths
  - Physical topology from router's TED or abstracted (=summarized) topology
- Out of scope:
  - Disseminating link-state information outside administrative domains

Transcoding TE Link Info into BGP NLRI

- Carried in
  - MP_REACH_NLRI
  - MP_UNREACH_NLRI
- Each NLRI describes a single link anchored by at least a pair of router-IDs
  - Link may be anchored by more than one pair of Router-IDs
- Negotiated between BGP speakers using BGP-MP Capability

Transcoding TE Link Info into BGP NLRI TED (SAFI 1)

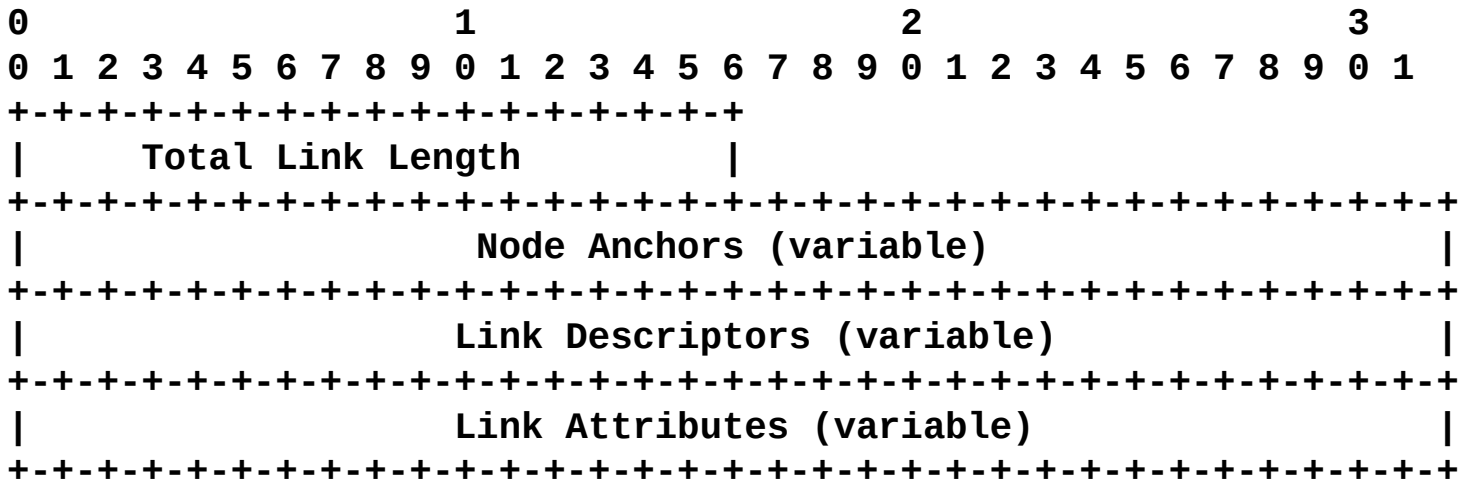

Node Anchor TLVs: describes which Protocols Router-IDs will "anchor" the link
Link Descriptor TLVs: uniquely identify a link between a pair of anchor Routers.
Link Attribute TLVs: describe the link properties

TED NLRI TLV Types

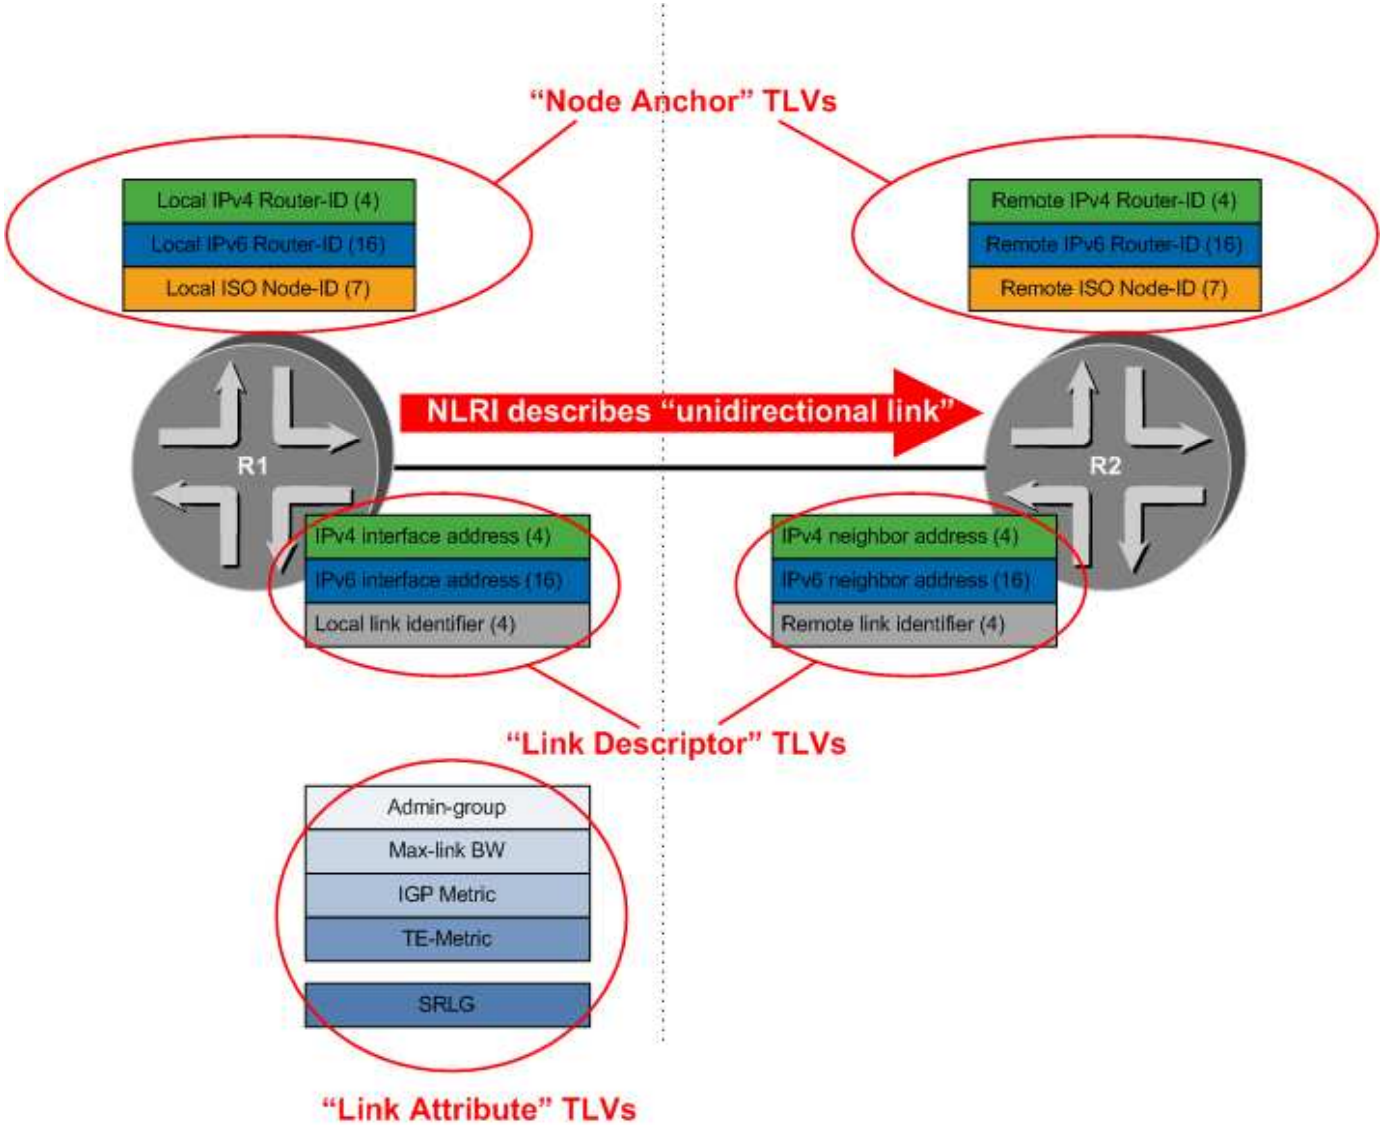

Transcoding TE Link Info into BGP NLRI TED (SAFI 128)

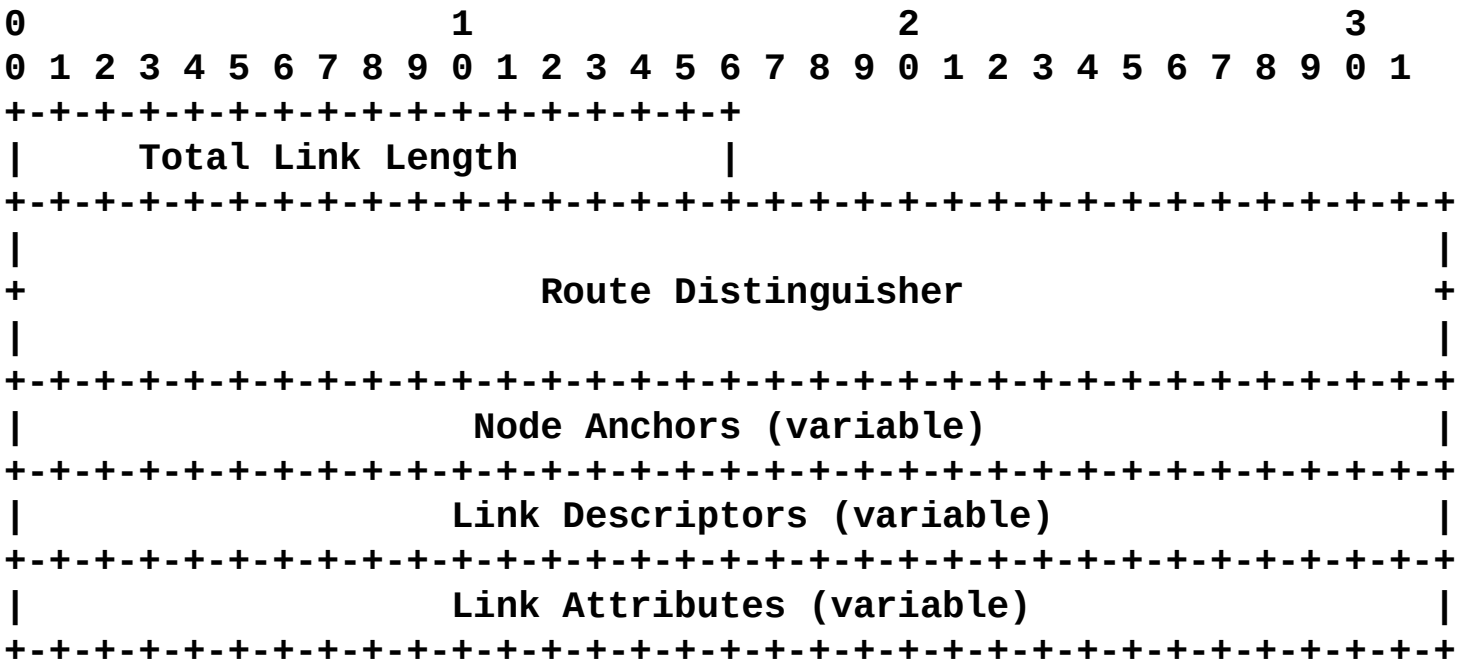

Route Distinguisher:
Node Anchor TLVs: describes which Protocols Router-IDs will "anchor" the link
Link Descriptor TLVs: uniquely identify a link between a pair of anchor Routers.
Link Attribute TLVs: describe the link properties

Node Anchors

Type	Description	Length
256	Local Autonomous System	4
257	Local IPv4 Router-ID	4
258	Local IPv6 Router-ID	16
259	Local ISO Node-ID	7
260	Remote Autonomous System	4
261	Remote IPv4 Router-ID	4
262	Remote IPv6 Router-ID	16
263	Remote ISO Node-ID	7

- **Local IPv4 Router ID:** opaque value (can be an IPv4 address or an 32 Bit router ID)
- **Remote IPv4 Router ID**: opaque value (can be an IPv4 address or 32 Bit router ID)
- **Local IPv6 Router ID**: opaque value (can be an IPv6 address or 128 Bit router ID)
- **Remote IPv6 Router ID**: opaque value (can be an IPv6 address or 128 Bit router ID)
- **Local ISO Node ID**: ISO node-ID (6 octets ISO system-ID plus PSN octet)
- **Remote ISO Node ID**: ISO node-ID (6 octets ISO system-ID plus PSN octet)
- **Local/Remote AS**: used to ***disambiguate*** Router-IDs allocated from private IP address spaces

Router-ID Anchoring Example ISO Pseudonode

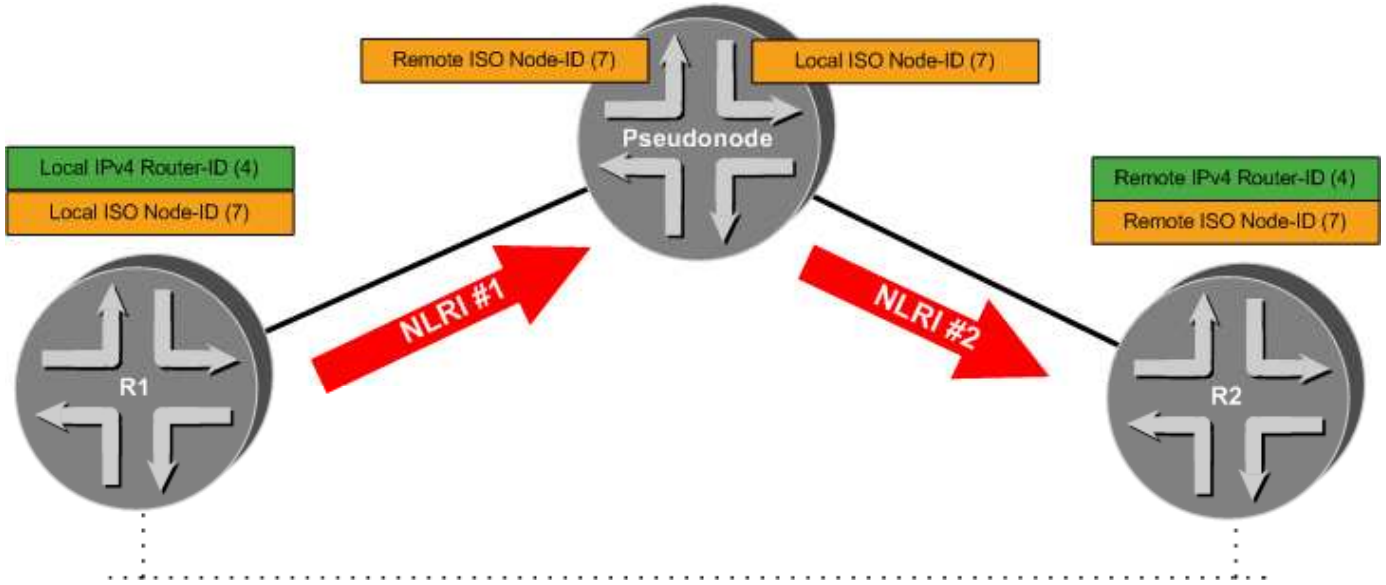

- Broadcast LAN between a pair of routers:
  - “Real" (=non pseudonode) routers have both an IPv4 Router-ID and IS-IS Node-IDs
  - The pseudonode does not have an IPv4 Router-ID.
- Two unidirectional links being generated:
- NRLI #1for (R1, Pseudonode) encodes:
  - local IPv4 router-ID, local ISO node-ID and remote ISO node-id
- NLRI #2 for (Pseudonode, R2) encodes:
  - local ISO node-ID, remote IPv4 router-ID and remote ISO node-id.

Link Descriptors

Type	Description	Defined in:
4	Link Local/Remote Identifiers	[RFC5307], Section 1.1
6	IPv4 interface address	[RFC5305], Section 3.2
8	IPv4 neighbor address	[RFC5305], Section 3.3
12	IPv6 interface address	[RFC6119], Section 4.2
13	IPv6 neighbor address	[RFC6119], Section 4.3

- Encoding of 'Link Descriptor' TLVs i(Type Codepoints, Lengths, Values) same as Extended IS reachability TLV sub-TLVs (defined in RFC5305, RFC5307 & RFC6119)
- Link Descriptor TLVs can carry data sourced either by IS-IS or OSPF.

Node Anchor + Link descriptor form the key in the DB/RIB

Link Attributes

Type	Description	Defined in:
3	Administrative group (color)	[RFC5305], Section 3.1
9	Maximum link bandwidth	[RFC5305], Section 3.3
10	Max. reservable link bandwidth	[RFC5305], Section 3.5
11	Unreserved bandwidth	[RFC5305], Section 3.6
20	Link Protection Type	[RFC5307], Section 1.2
64512	TE Default Metric	Section 3.4.1
64513	IGP Link Metric	Section 3.4.2
64514	Shared Risk Link Group	Section 3.4.3

- Where possible, TLV definitions from IS-IS TE reused:
  - For Codepoints < 255 Encoding of 'Link Descriptor' TLVs i(Type Codepoints, Lengths, Values) same as the ‘Extended IS reachability TLV’ sub-TLVs
- New TLVs where link attributes encoded differently in OSPF-TE and IS-IS TE:
  - TE Default Metric
  - IGP Link Metric
- The SRLG TLV:
  - Same info as in IS-IS SRLG TLVs (Type 138, Type 139)
  - No need to include Router-ID link anchors in the TLV

Next steps

- Publish draft-gredler-bgp-te-01.txt
  - Clarify NLRI key (link anchors plus link attributes)
  - Describe path selection
  - Clarify Security section
  - Add support for Node Properties (Overload bit etc).