



Segment Routing Prefix SID extensions for BGP *draft-keyupate-idr-bgp-prefix-sid-05*

Stefano Previdi (sprevidi@cisco.com)

Clarence Filfsils (cfilfil@cisco.com)

Keyur Patel (keyupate@cisco.com)

Arjun Sreekantiah (asreekan@cisco.com)

Saikat Ray (raysaikat@gmail.com)

Hannes Gredler (hannes@juniper.net)


Acee Lindem (acee@cisco.com)

BGP Prefix-SID Attribute

- New BGP Attribute addressing the use case described in draft-filsfils-spring-segment-routing-msdc
- Version 5
 - Merged with draft-gredler-idr-bgplu-prefix-sid-00
 - Added support of SR-IPv6 dataplane with SR-IPv6-SID TLV
 - Added Originator SRGB TLV when SRGB is to be learned through Prefix-SID attribute
 - Applicable to Labeled unicast prefixes (RFC3107) and MP-BGP unlabeled unicast IPv6 prefixes (RFC4760)
- Multiple implementations exist
- Ready for WG adoption

BGP Prefix-SID Attribute

- Label Index TLV




where:


- Type is 1.
- Length: is 7, the total length of the value portion of the TLV.
- RESERVED: 8 bit field. SHOULD be 0 on transmission and MUST be ignored on reception.
- Flags: 16 bits of flags. None are defined at this stage of the document. The flag field SHOULD be clear on transmission and MUST be ignored at reception.
- Label Index: 32 bit value representing the index value in the SRGB space.

BGP Prefix-SID Attribute

- SR IPv6 SID



- o Type is 2.
- o Length: is 3, the total length of the value portion of the TLV.
- o RESERVED: 8 bit field. SHOULD be 0 on transmission and MUST be ignored on reception.
- o Flags: 16 bits of flags defined as follow:




where:

S flag: if set then it means that the BGP speaker attaching the Prefix-SID Attribute to a prefix is capable of processing the IPv6 Segment Routing Header (SRH, draft-previdi-6man-segment-routing-header) for the segment corresponding to the originated IPv6 prefix.

BGP Prefix-SID Attribute


- Originator SRGB



- o Type is 3.
- o Length is the total length of the value portion of the TLV: 2 + multiple of 6.
- o Flags: 16 bits of flags. None are defined in this document. Flags SHOULD be clear on transmission and MUST be ignored at reception.
- o SRGB: 3 octets of base followed by 3 octets of range. Note that SRGB field MAY appear multiple times.

BGP Prefix-SID Attribute

- Originator SRGB TLV



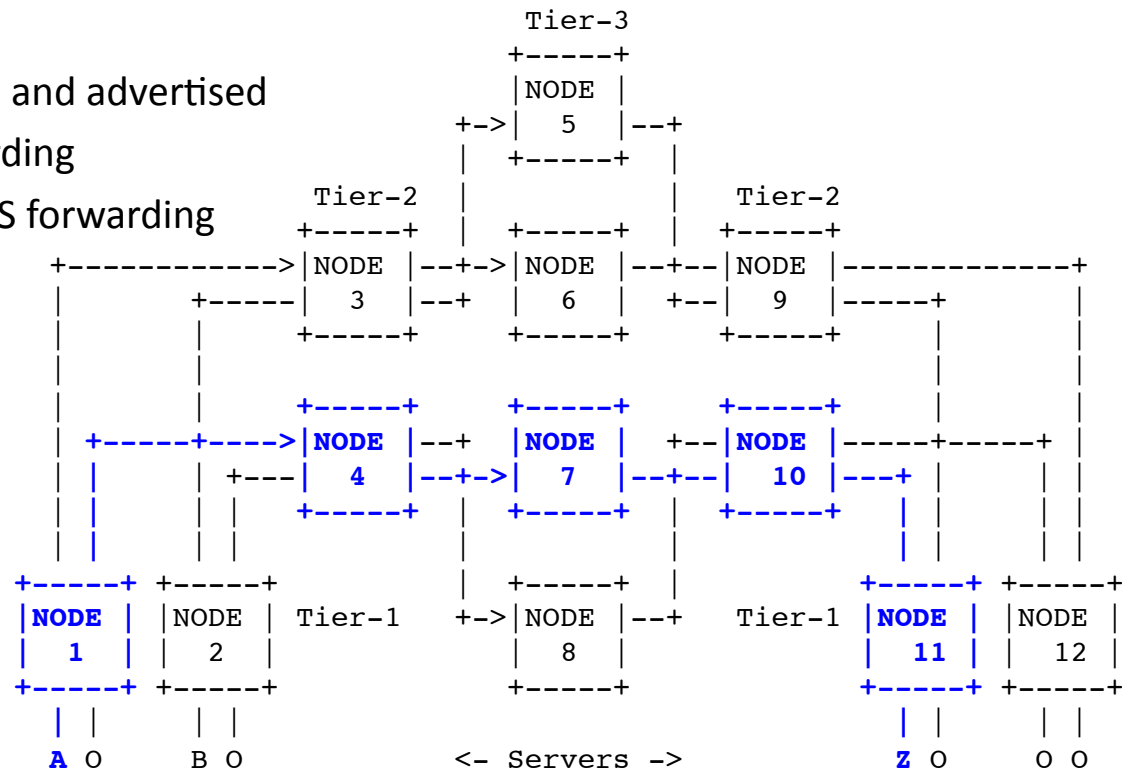
The Originator SRGB TLV contains the SRGB of the router originating the prefix to which the BGP Prefix SID is attached and MUST be kept in the Prefix-SID Attribute unchanged during the propagation of the BGP update.

The originator SRGB describes the SRGB of the node where the BGP Prefix Segment end. It is used to build SRTE policies when different SRGB's are used in the fabric (draft-filsfils-spring-segment-routing-msdc).

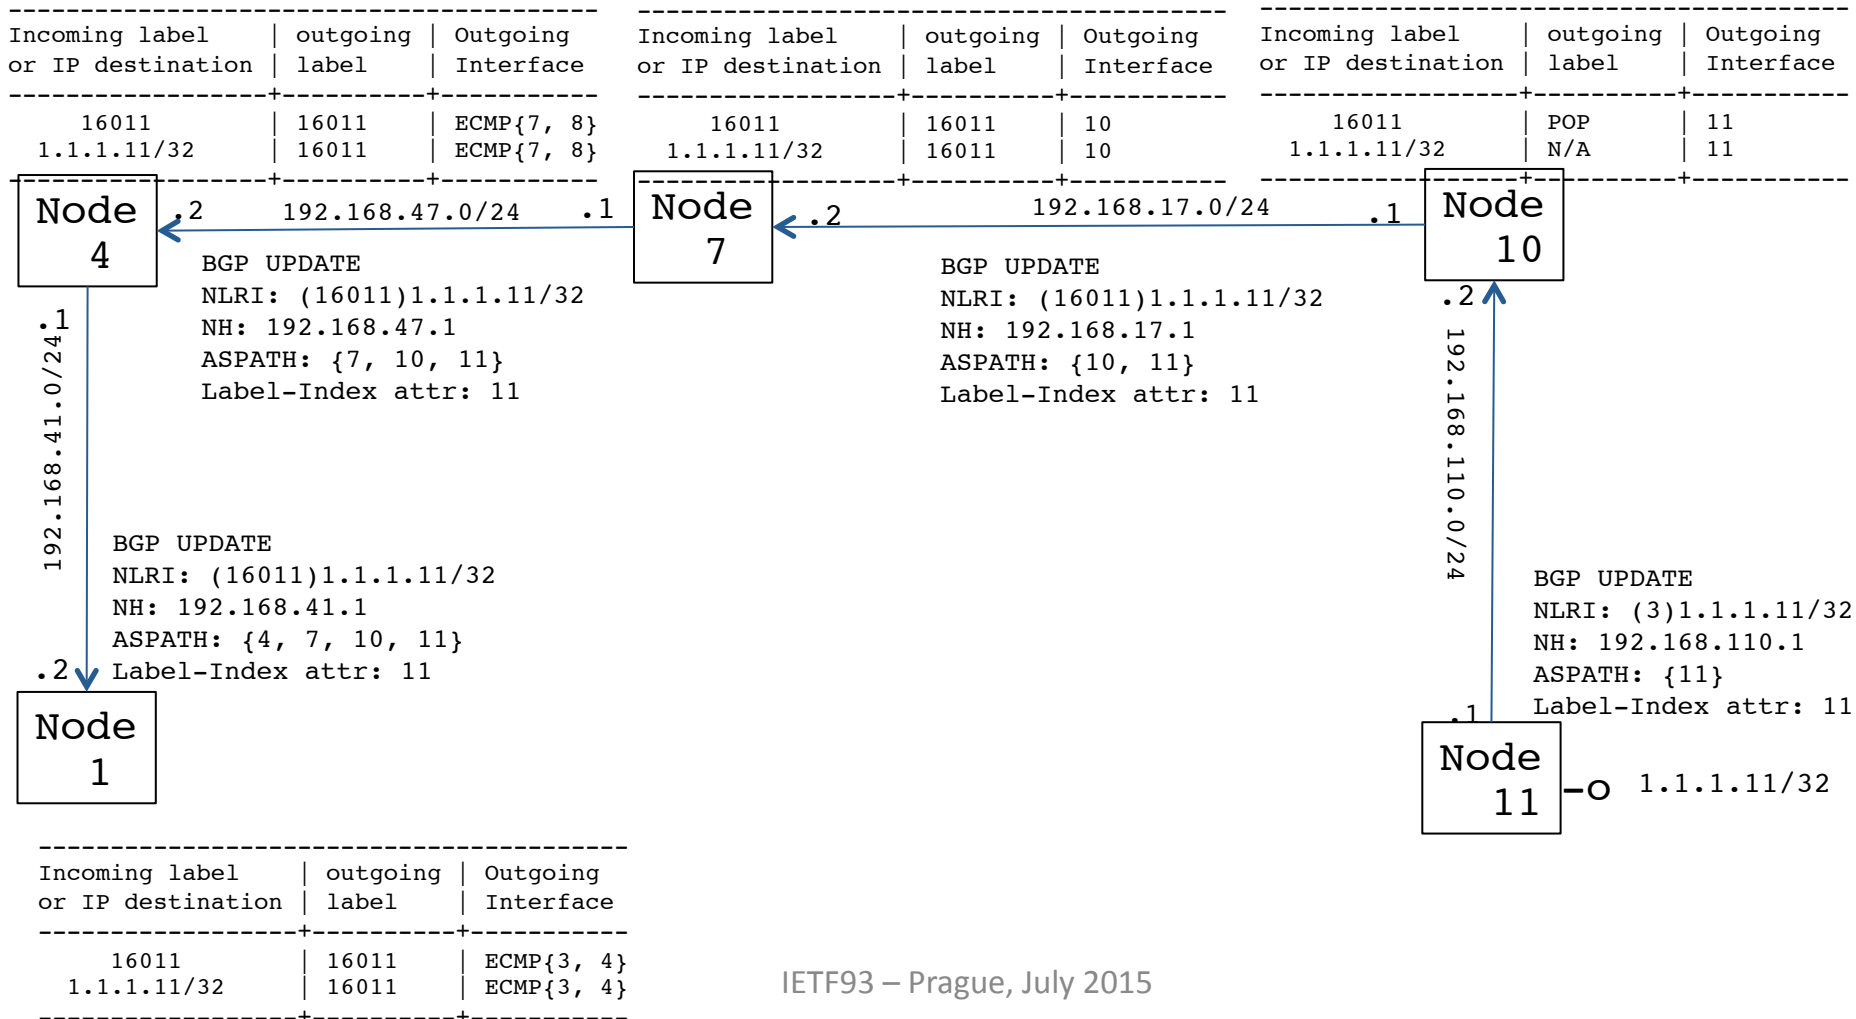
The originator SRGB may only appear on Prefix-SID attribute attached to prefixes of SAFI 4 (labeled unicast, [RFC3107]).

Reference topology

- Node 'x' has ASN 'x'
- BGP IPvX labeled-unicast sessions (3107) between directly connected nodes
- Node 'x' has loopback 1.1.1.x/32
- Loopbacks are redistributed into BGP and advertised
- Tier-2 and Tier-3 nodes: MPLS forwarding
- Tier-1 nodes: IP2MPLS or MPLS2MPLS forwarding
- SRGB: [16000, 23999]
- Label index for 1.1.1.x/32 is 'x'



BGP Prefix SID: Control and dataplane



BGP Prefix SID: Non-SR node in the middle

