IETF 111 – Online
July 2021

BGP Color-Aware Routing (CAR)

draft-dskc-bess-bgp-car-02

D. Rao, Cisco Systems (dhrao@cisco.com) – Presenter
S. Agrawal, Cisco Systems
C. Filsfils, Cisco Systems
K. Talaulikar, Cisco Systems
L. Jalil, Verizon
Y. Su, Alibaba
D. Steinberg, Steinberg Consulting
K. Patel, Arrcus
H. Wang, Huawei
J. Guichard, Futurewei

BGP Color-Aware Routing

- Define BGP based routing solution to establish end-to-end intent-aware paths across a multi-domain network environment
  - Intent : Example – low-latency path between two PEs

- Color represents intent in signaling
  - draft-ietf-spring-segment-routing-policy
  - draft-ietf-idr-segment-routing-te-policy

BGP Color-Aware Route & Automated Steering

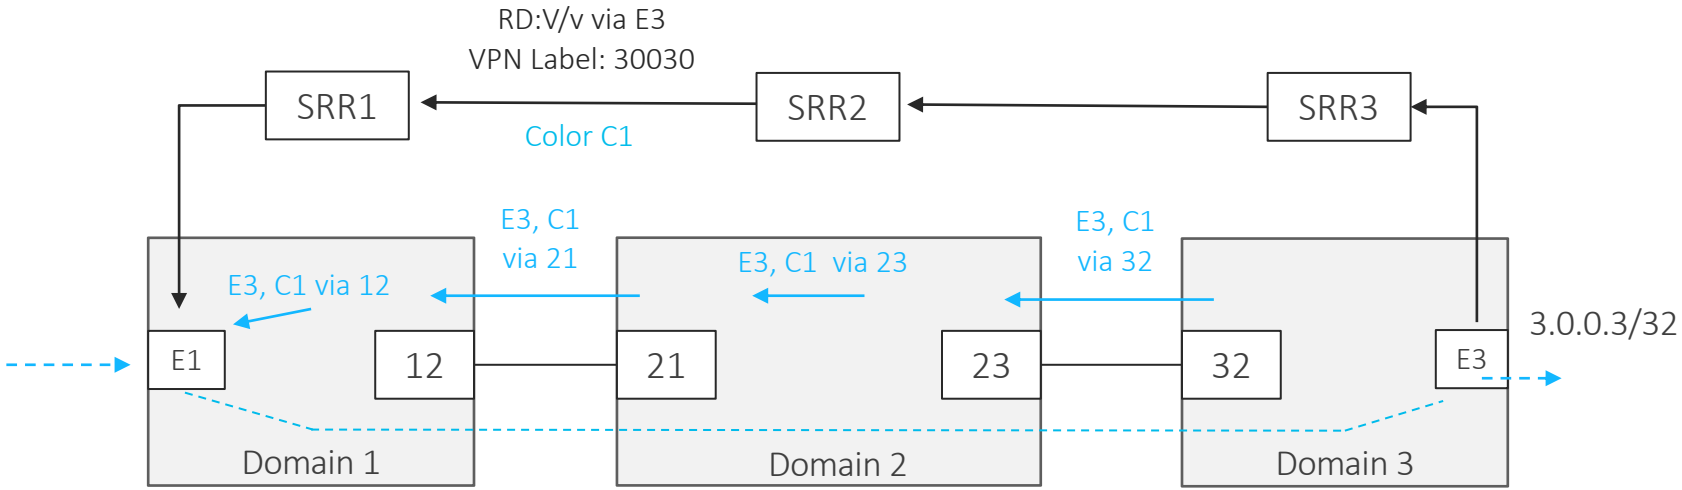

- E3, C1 is a Color-Aware BGP route in underlay that provides intent-aware path to E3

BGP CAR Overview (Refresher)

- New SAFI in BGP
  - Need ability to signal multiple instances of the same prefix for each color (i.e., intent)
- Solution draft v01 described the following aspects
  - Desired Data Model
  - Multiple encapsulations, their signaling and validation
  - Efficient and extensible NLRI
  - Handling of multiple color domains
  - Route resolution & steering mechanisms
  - Scale Analysis

CAR NLRI Proposal

- NLRI Key – E, C
  - E : IPv4 or IPv6 Endpoint Prefix (Network-wide Unique)
  - Color : 32-bit value (same as SR-TE Policy)

- Color distinguishes per-intent instances of same prefix
- Color also indicates intent provided by route
- Color is consistent across devices within a “color domain”
- Color is same as in BGP Color Extended-Community

CAR NLRI – E, C

- Simplest data model, precise
- Identical routing semantics as BGP IPv4/v6, BGP-LU
  - Efficient route processing, storage
  - No need for VPN import/export or RD rewrite at each underlay hop
- Inherently provides ECMP-aware/backup paths at every hop
  - Faster, localized convergence
- Most efficient for subscription
  - [E, C] direct lookup

Path Availability & Domain-local Convergence

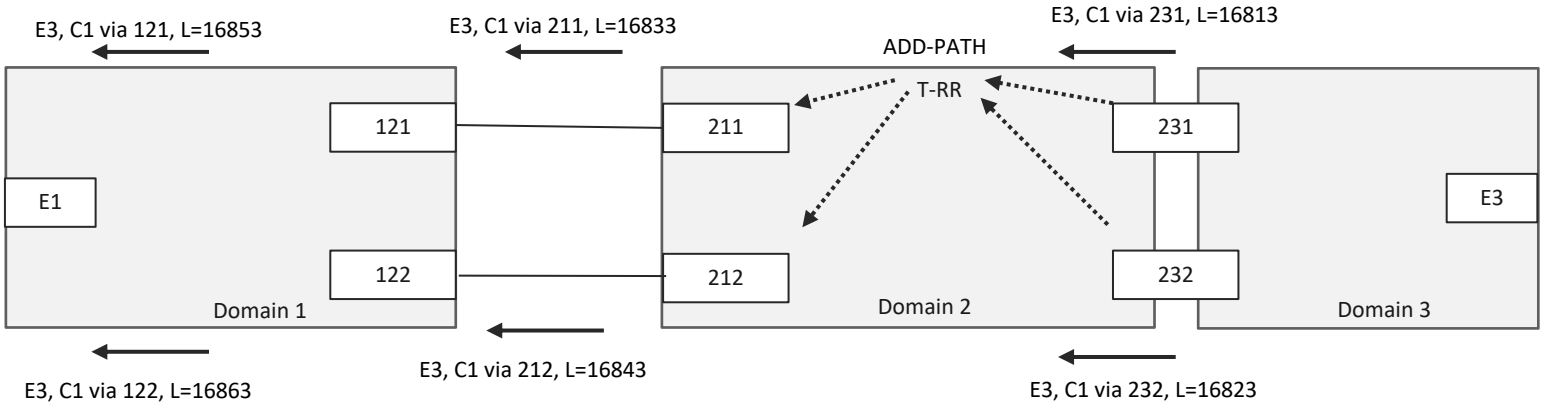

- (E, C) NLRI provides ECMP or backup paths at each hop (single label entry)
- Localized convergence with Next-Hop Self
  - E.g., 231 failure is handled locally within domain, churn is not propagated beyond 212 and 211
- BGP ADD-PATH at T-RR for redundant path availability

Multiple Color Domains

- Network domains where color-intent mappings are different
- Local-Color-Mapping (LCM) Extended Community
  - Optional, only used if routes go across a color domain boundary
  - Color re-mapped and rewritten into receiving domain's color at a color domain boundary
  - Color Ext-Comm sent with service routes also gets re-mapped in parallel
- CAR NLRI (E, C) is preserved e2e
- E (Prefix) is unique in inter-domain transport network (e.g., PE)
  - Makes E, C unique even if C is local to a color domain

Encapsulations

- Multiple encapsulations supported for a CAR route
  - Signaled via Non-Key TLVs
    - MPLS Label(s), Label-Index, SRv6 SID(s) etc.
  - Separate “label” values for different encapsulations
  - Beneficial for co-existence, migration & interworking
    - Efficient signaling, operational simplicity

Extensible, Future-Proof NLRI Encoding

- New SAFI allows opportunity for better NLRI design
  - No need to inherit constraints of current SAFIs, e.g., single MPLS label field in NLRI
- Encode a NLRI (Route) Type
- Encode a key length
- Encode non-key TLVs
- Variable part in NLRI; rest in Attribute
  - Provides packing efficiency for BGP updates

CAR Next-Hop Resolution

- Resolution is recursive and color-aware
  - (E, C) via (N, C)
- (N, C) provided by other color-aware mechanisms
  - SR Policy, IGP Flex-Algo, or BGP CAR itself
- Resolution may also be mapped to legacy mechanisms
  - RSVP-TE
  - IGP/LDP
  - BGP-LU

Updates (v02)

- VPN CAR - Extends CAR to VPN service layer
  - CE – PE BGP Color-aware routing
  - E2E encapsulation (e.g., CE – CE)
- RFC 4364 semantics
  - VPN RD, RT(s), Import/Export
- CAR NLRI requires new SAFI
  - Straightforward extension of CAR NLRI - (RD : E, C)
    - Where RD is regular VPN RD
  - No overloading of RD for both VPN & Color separation

VPN CAR Illustration (Simplified)

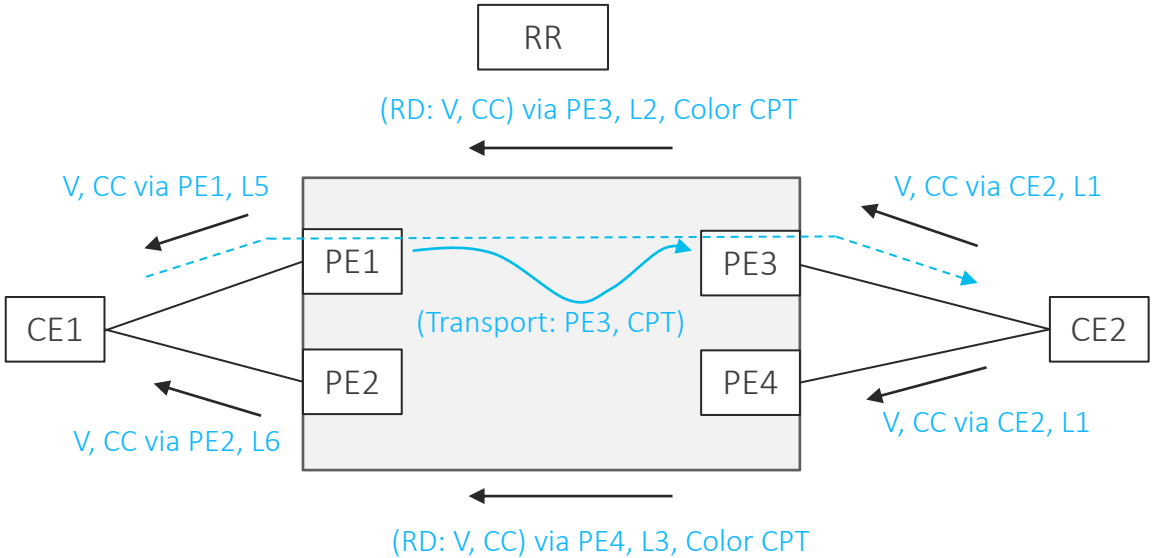

Updates (v02)

- Describe usage of Anycast SID
  - Convergence, Recursion
  - Anycast SID for transit, e.g., ABRs
  - Anycast SID for PEs (with common service labels)

- Clarify path availability & convergence
  - Covered in earlier slide

- Added J. Guichard as co-author

Next Steps

- Continue to address use-cases & requirements listed in problem statement
- Request collaboration & review from Working Group
- Problem statement drafts merge effort is ongoing