IETF 110 – Online
March 2021

BGP Color-Aware Routing (CAR) Problem Statement

draft-dskc-bess-bgp-car-problem-statement

D. Rao, Cisco Systems (dhrao@cisco.com) – Presenter
S. Agrawal, Cisco Systems
C. Filsfils, Cisco Systems
K. Talaulikar, Cisco Systems
B. Decraene, Orange
D. Steinberg, Steinberg Consulting
L. Jalil, Verizon
J. Guichard, Futurewei
W. Henderickx, Nokia

BGP Color-Aware Routing - Objective

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

Reminder – Deployed Solution

- SR-TE
  - ietf-spring-segment-routing-policy
  - Mature, widely deployed, multiple implementations
  - Defines notion of Color to represent intent

Colored Service Route Signaling from E3 to E1

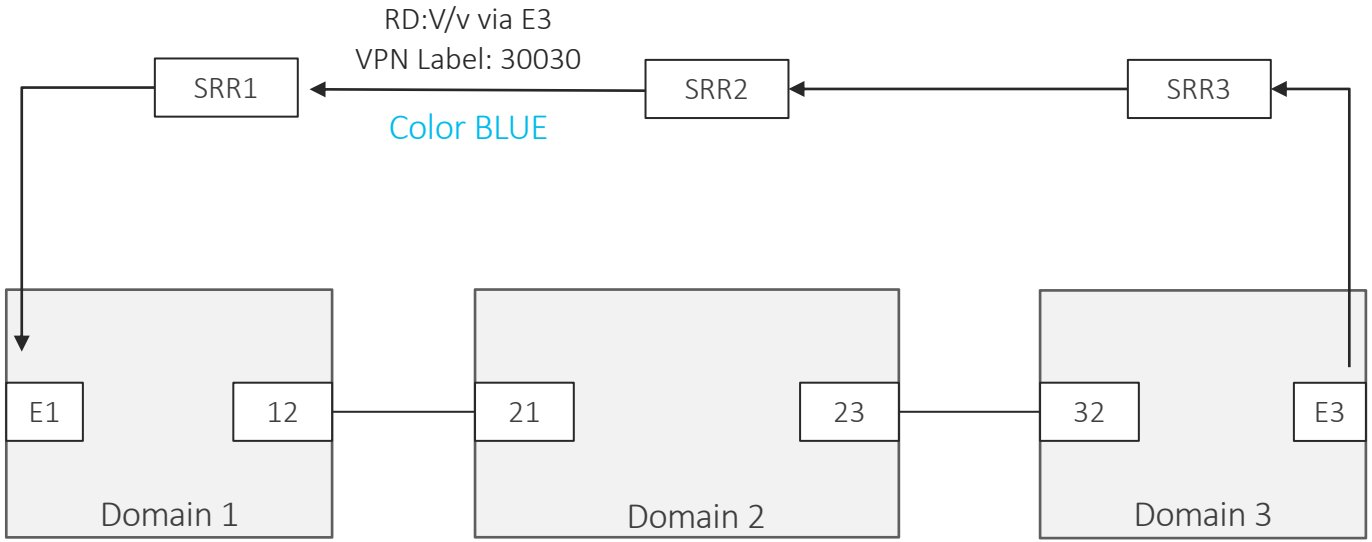

- Key point: E1 learns about the “intent” (here for underlay SLA) requested by a route via its color
- The VPN route is said to be “colored” (<> color-aware)
- Color is widely supported BGP Color Extended-Community

Automated Steering via SR-TE Color-Aware Path

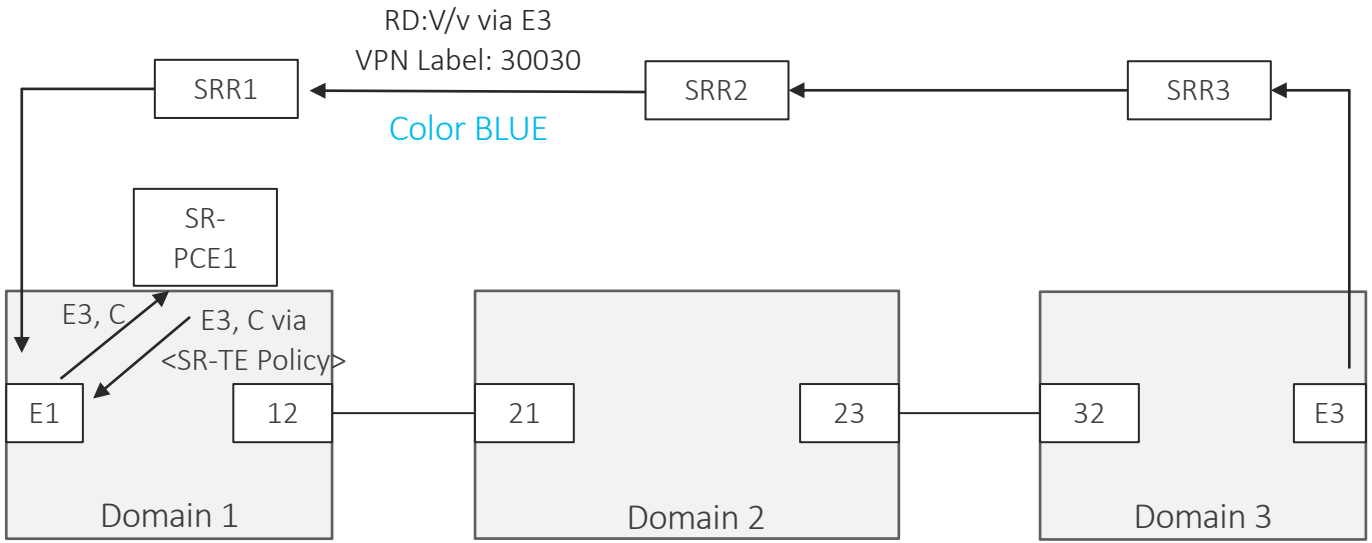

- When E1 receives a Colored Service route from E3
- E1 requests its SR-PCE1 to compute the inter-domain path
- SR-PCE1 sends the SR Policy to E1 with label/SID stack
- E3, C is a Color-Aware Path in underlay that provides intent-based path to E3

Automated Steering Evolution - BGP Color-Aware Route

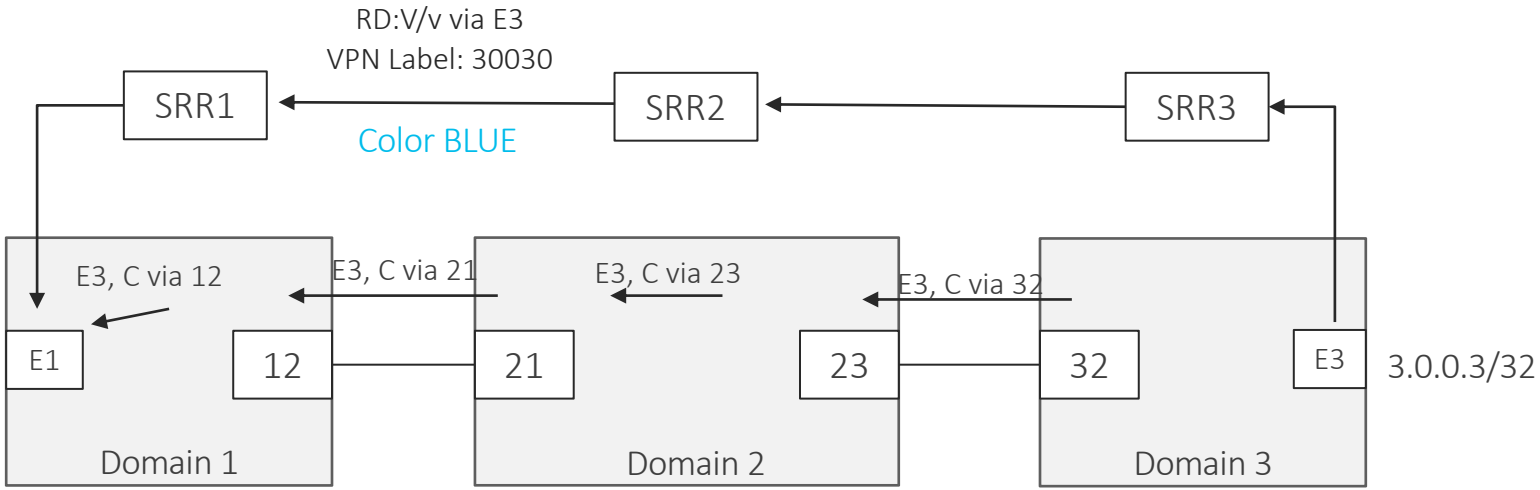

- E3, C is a Color-Aware BGP route in underlay that provides intent-aware path to E3

Focus of Problem Statement Draft

- Clarity on deployment requirements
  - Compatibility, co-existence, interworking with deployed SR-PCE based solution
  - Color to drive automated steering
- Widened problem scope
  - Intent-aware VPN service layer
  - NFV Integration
- Clarity on Scale requirements and constraints
  - Data Plane (MPLS label space / FIB)
  - Control Plane (BGP) Filtering
- Crisp, technical analysis of intent use-cases and protocol requirements

Collaboration

- Collaboration with lead operators, vendors on analysis
  - Acknowledge many contributors in draft
- Recognize prior work
  - Seamless SR/Classful Transport
- Ongoing collaboration effort with SSR co-authors for consensus
  - Reached out through co-authors in Nov/Dec
  - Recognized prior publication on use-cases / illustrations
  - We published problem statement with analytical approach as contribution
  - SSR co-authors acknowledged feedback & split their document
  - Joint discussion progressing well for eventual partnership, new sets of documents

Next Steps

- Request review from Working Group

IETF 110 – Online
March 2021

BGP Color-Aware Routing (CAR)

draft-dskc-bess-bgp-car

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

New SAFI for BGP CAR

- Need for a new SAFI in BGP
  - Need ability to signal multiple instances of the same prefix for each color (i.e. intent)
- Solution draft describes the following aspects
  - Desired Data Model
  - Multiple encapsulations, their signaling and validation
  - Route resolution & steering mechanisms
  - Scale Characterization
  - Route Filtering

Data Model

- Prefix : IPv4/IPv6 (global scope)
- Color : Distinguishes different instances of a Prefix, per intent

  : Indicates the service/intent associated with the route
- Next Hop : as usual; reachability verified via routing control plane
- Encapsulation :
  - Support for state-of-art : MPLS label(s), SRv6 SID(s), IP etc
  - Support multiple encapsulations
  - AIGP Metric : for accumulation of intent-specific metric

CAR NLRI Proposal

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

Multiple Color Domains

- Local-Color-Mapping (LCM) Extended Community
  - Optional, used only if routes go across a color domain boundary
  - Color re-mapped and rewritten into receiving domain's color at a color domain boundary
- CAR NLRI (E, C) is preserved e2e

Encapsulation

- Variable part in NLRI; rest in Attribute
  - Necessary for packing efficiency of BGP updates
  - Opportunity for clean design that is not constrained by 24-bit MPLS label field in the NLRI
- Validate data-plane availability of encapsulation before using

Scale Considerations

- Hierarchical Design
  - Hierarchical models avoid the need for core BRs to learn routes and install label forwarding entries for (E, C) routes
  - Analyze recursion and data plane complexity at ingress PE/BR
- Filtering
  - Ingress PE/BR only learns (E, C) routes that it needs to install into data plane

Next Steps

- Request review from Working Group