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 path between two PEs

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

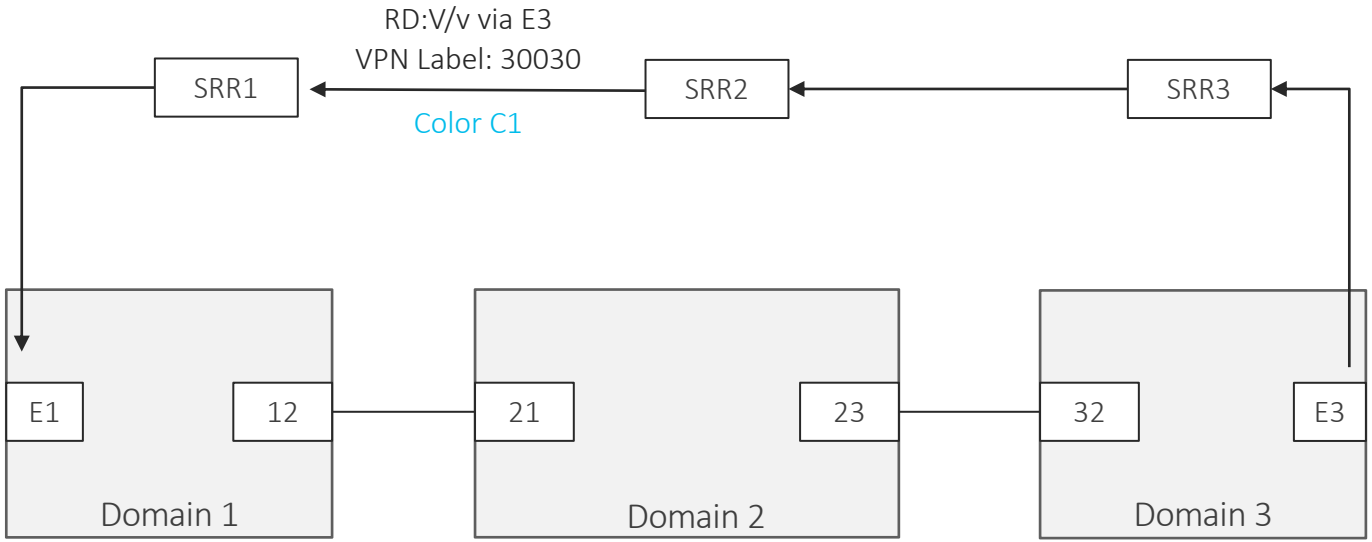

- Key point: E1 learns about the “intent” (here for underlay SLA) requested by a route via its color
- The VPN route is said to be “colored” (<> color-aware)
- Color is widely supported BGP Color Extended-Community

Automated Steering via SR-TE Color-Aware Path

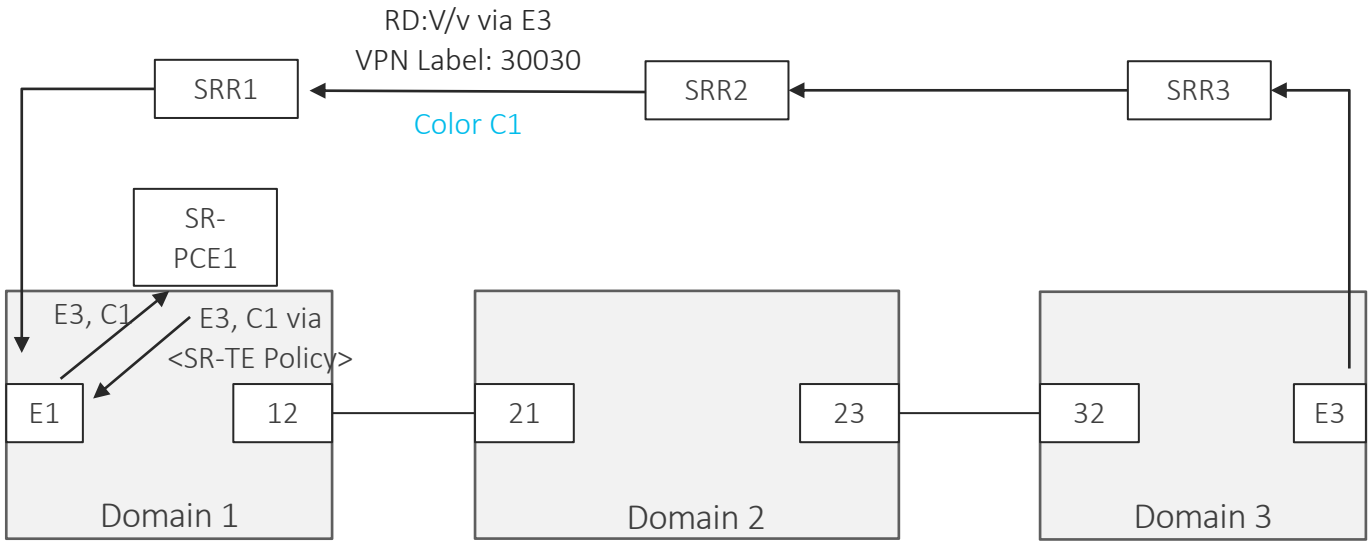

- When E1 receives a Colored Service route from E3
- E1 requests its SR-PCE1 to compute the inter-domain path
- SR-PCE1 sends the SR Policy to E1 with label/SID stack
- E3, C1 is a SR-Policy Color-Aware Path in underlay that provides intent-aware path to E3

Automated Steering Evolution - BGP Color-Aware Route

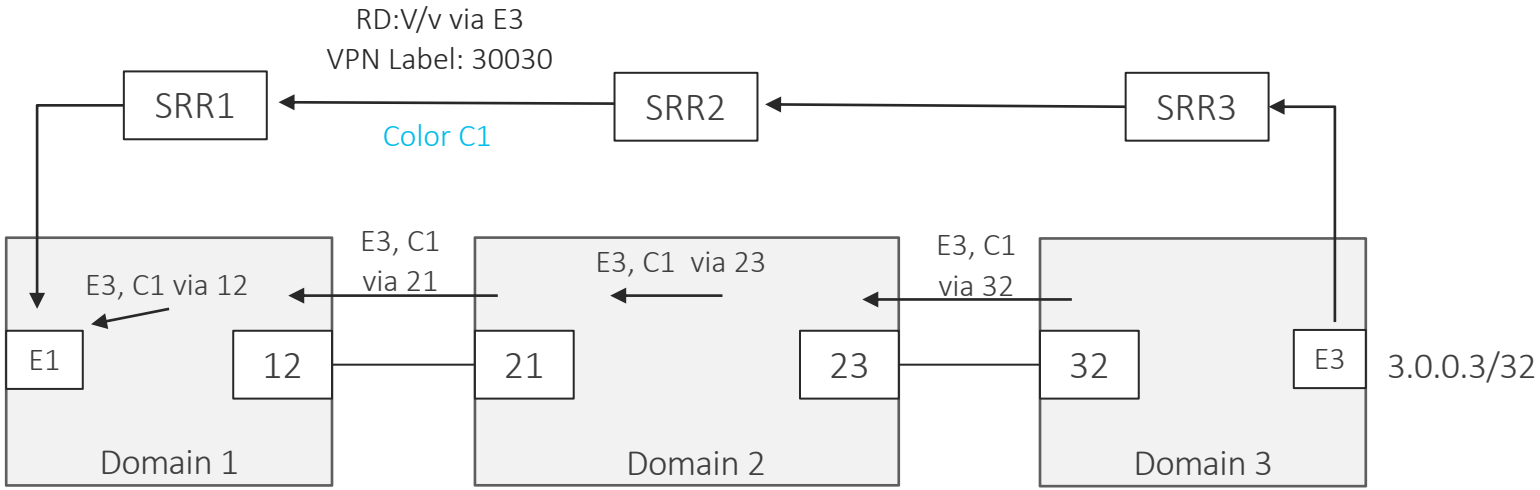

- E3, C1 is a Color-Aware BGP route in underlay that provides intent-aware path to E3

Multiple Intents with BGP Color-Aware Routes

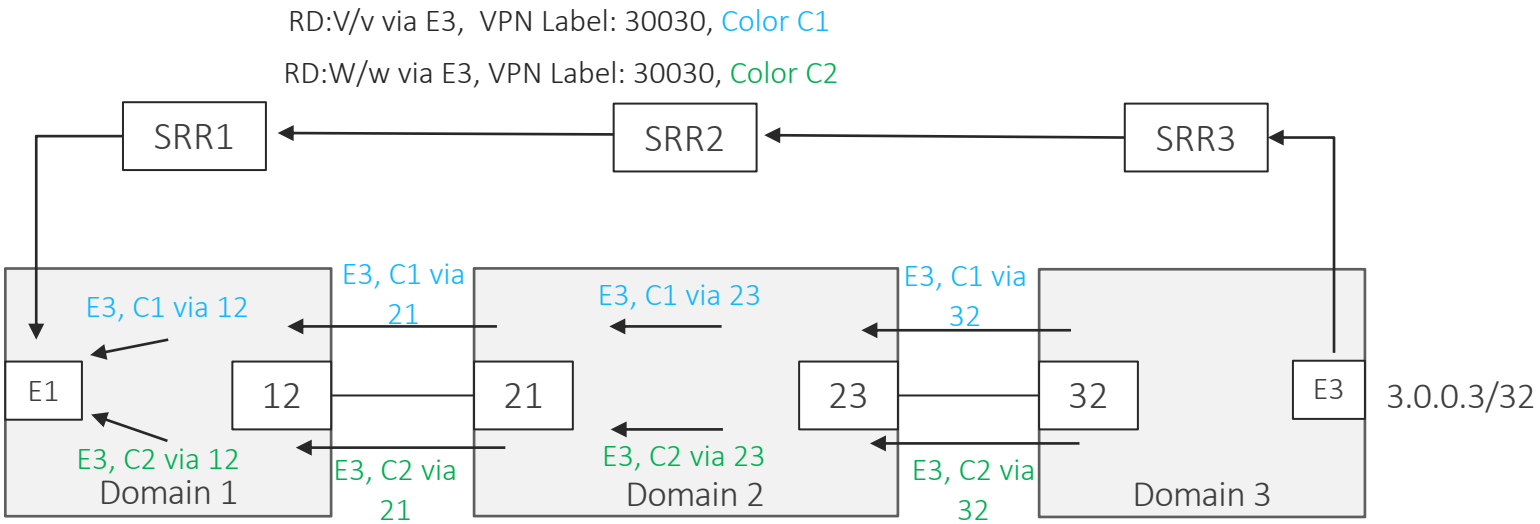

Reference Deployment Design

- Well-known MPLS network reference deployment designs:
  - Seamless MPLS
  - Inter-AS option C
- Ultra-large-scale multi-domain network with around 300K nodes
  - Core, Metro, Aggregation, Access layers
- Multiple intents (1 best-effort and 4 intents for example)
  - Low-latency
  - Plane 1 & Plane 2
  - Avoidance (links/nodes/domains – for regulatory, security, quality, etc.)

Types of Intent bound to a Color

- Minimization of different metrics – link cost, latency
  - Minimization of different metric types, static and dynamic
- Exclusion/Inclusion of SRLG and/or Link Affinity
- In the inter-domain context, exclusion/inclusion of entire domains, and border routers
- Minimum MTU / number of hops / MSD
- Bandwidth management, to the extent possible
- Inclusion of one or several virtual network function chains
  - Localization of the virtual network function chains

Focus of Problem Statement Draft

- Crisp, technical analysis of intent use-cases and protocol requirements
- Consistency, co-existence, interworking with deployed SR-Policy based solution
  - Color to drive automated steering
- Widened problem scope
  - Intent-aware VPN service layer
  - NFV Integration

Problem Statement Draft Contd.

- Clarity on deployment requirements
  - E2E paths across domains with different technologies and encapsulations

- Clarity on Scale requirements and constraints
  - Data Plane (MPLS label space / FIB)
  - Control Plane (BGP) Filtering

Collaboration

- Collaboration & review 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
  - Efficient and extensible NLRI
  - Route resolution & steering mechanisms
  - Scale Analysis
  - Route Filtering

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 color-aware semantics as SR-Policy
- Similar routing semantics as BGP IPv4/v6, BGP-LU
  - Efficient route processing, storage
- Inherently provides ECMP-aware/backup paths at every hop
  - Faster, localized convergence
- Most efficient for subscription
  - [E, C] direct lookup

Encapsulation

- Multiple encapsulations for a route
  - Non-Key TLVs
    - MPLS Label(s), Label-Index, SRv6 SID(s) etc
  - Co-existence, migration, interworking

- Variable part in NLRI; rest in Attribute
  - Necessary for packing efficiency of BGP updates
  - Opportunity for clean design not constrained by 24-bit MPLS label field in NLRI

Extensible, Future-Proof NLRI Encoding

- Encode a NLRI (Route) Type
- Encode a key length
- Encode non-key TLVs

AIGP

- Re-use to carry accumulated metric for specific color (intent)
  - Extend for intent-specific requirements
  - Indicate discontinuity in intent

CAR Route Validation & Resolution

- Validation
  - Availability of Color-Aware Path to Next-Hop (e.g., an IGP FA)
    - Policy to fallback to default
  - Data-plane availability of encapsulation
  - Performance measurement
- Resolution
  - Recursive via BGP or intra-domain
    - SR Policies, IGP Flex-Algo
    - IGP best effort (SR, LDP, RSVP-TE, etc.)
    - BGP CAR / BGP-LU

Service Route Automated Steering

- As seen in problem statement
  - Via BGP CAR, SR-Policy, IGP-FA
  - Flexible for all encapsulations

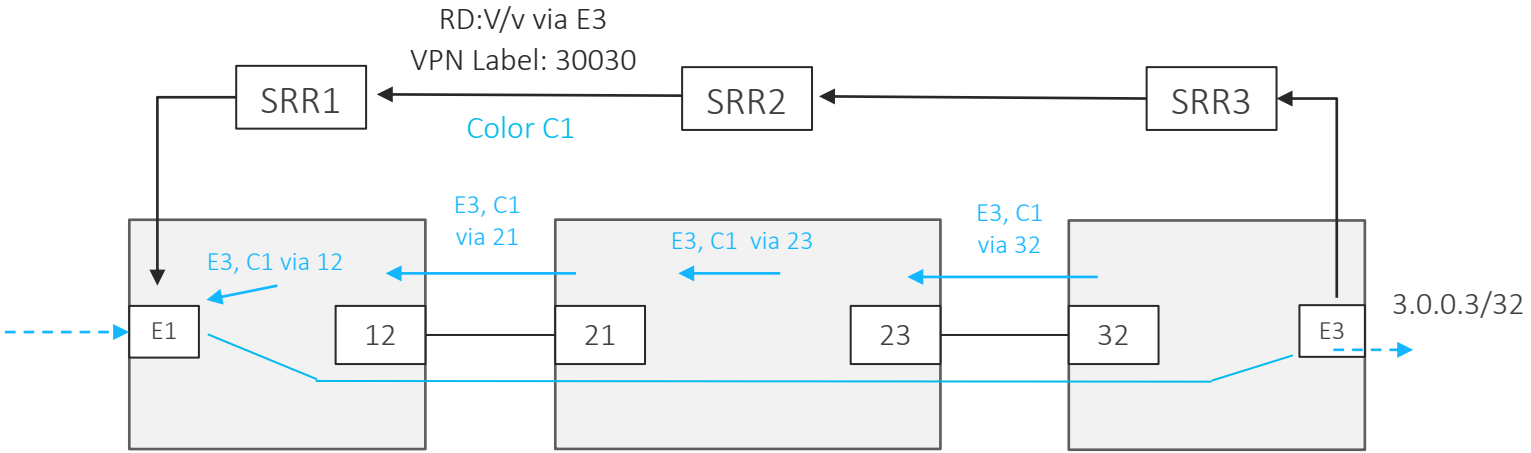

Multiple Color Domains

- 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
- 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

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