

CAR-CT Adoption Call

7/6 to 7/27

IDR chairs: Susan Hares, Jeff Haas, Keyur Patel

Dhananjaya Rao, Swadesh Agrawal

Kaliraj Vairavakkalai, Natrajan Venkataraman

Adoption call (7/6 to 7/27)

- 48 people responded
 - Part 1 (20) – Customer Intent is SLA+SLE, Need interoperable, scalable solution
 - Part 2 (48) – Support
 - Part 3 (19) - [Number of issues raised]
- Input from operators, customer support, implementers, academics
- Operators and customers determined course

Daniel Bernier (Bell Canada), Gaurav Dawra (Linkedin), Bruno Decraene (Orange), Amit Dhamija (Rakuten), Goekhan Guemues (eunetworks), Luay Jalil (Verizon), Satoru Matsushima, Israel Means (AT&T), Gyan Mishra (Verizon), Moses Nagarajah (Telestra), Robert Raszuk, Yuancaho Su (Alibaba), Tomasz Szewczyk (PSNC) , Jim Uttaro (AT&T), Aijun Wang (China Telcom), and Shawn Zhang.
- IDR Wiki has Summaries and notes

Adoption Call Result

- 2 Drafts adopted as Experimental track
 - draft-dskc-bess-bgp-car-05
 - draft-kaliraj-idr-bgp-classful-transport-planes-17
 - WG generates CAR-CT Interoperability document
- Why Experimental – CAR and CT
 - Customers convinced chairs that technology is needed in field
 - Interoperability between CAR-CT – less important than deploying code
 - WG needs to turn Drafts → RFC quality
- Experimental to Proposed standard
 - Revise drafts from customer feedback
 - Create Interoperability Draft
 - 2 implementations

Topics

Customer Examples:

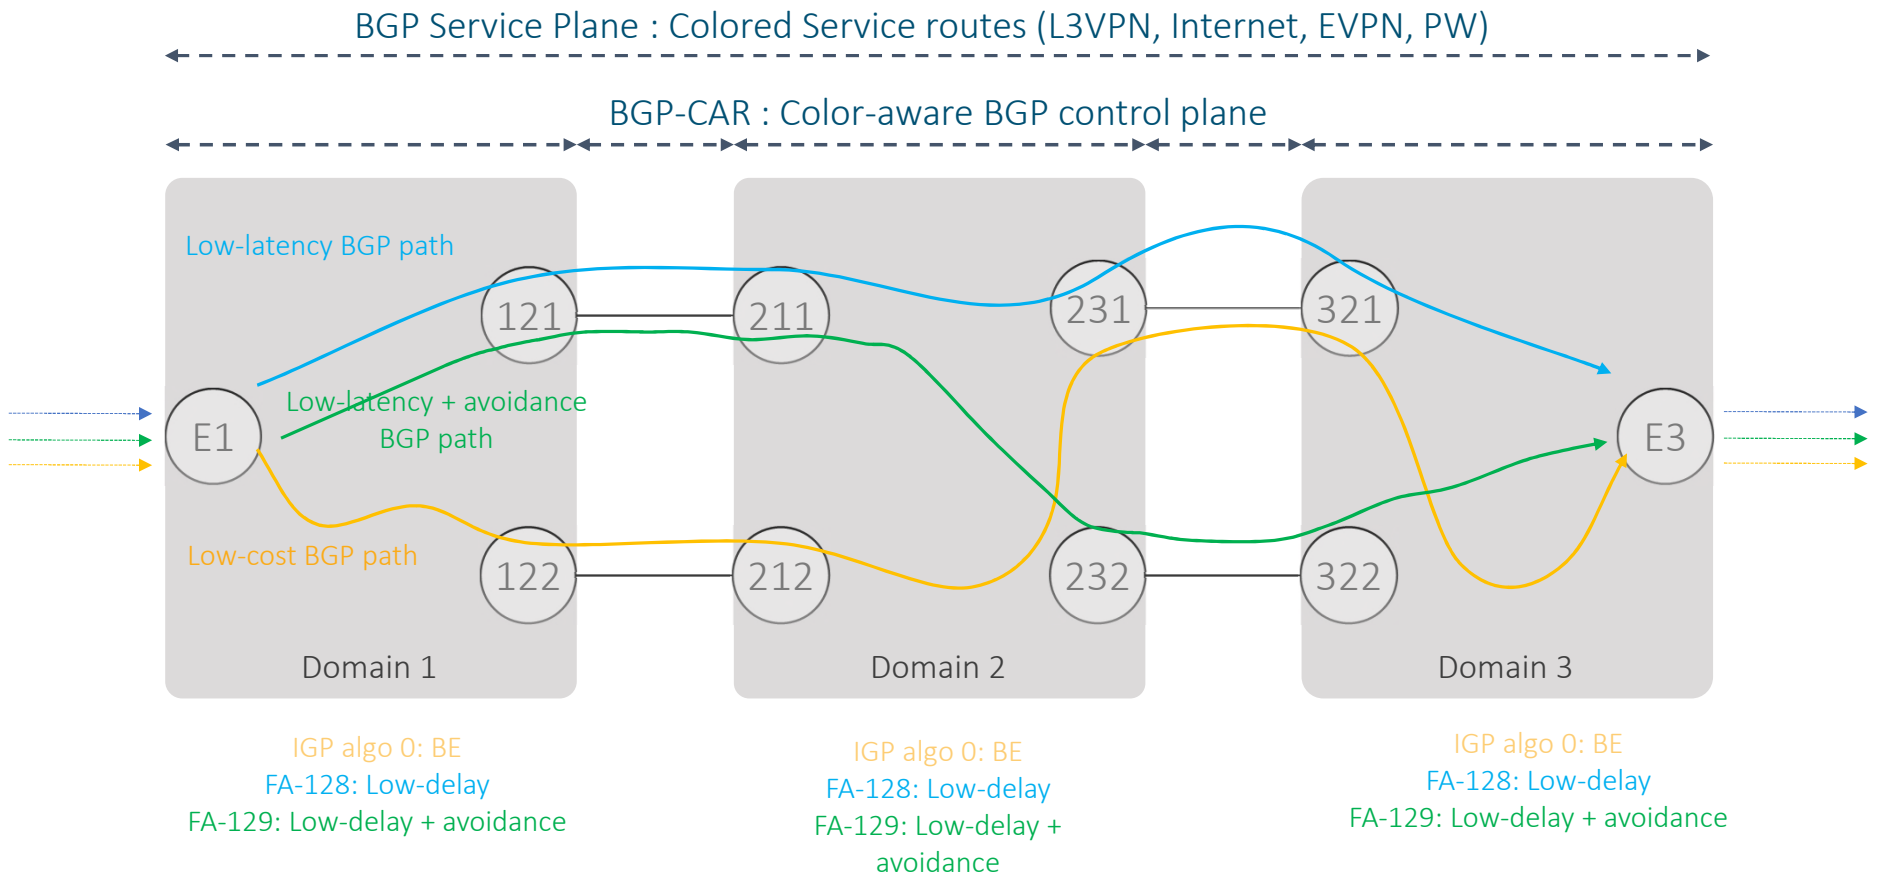
- CAR customer's
- CT customer's

Next Steps:

- Adoption Call – Part 3 issues
- CAR Next Steps
- CT Next Steps
- Interoperability Document Starting Point

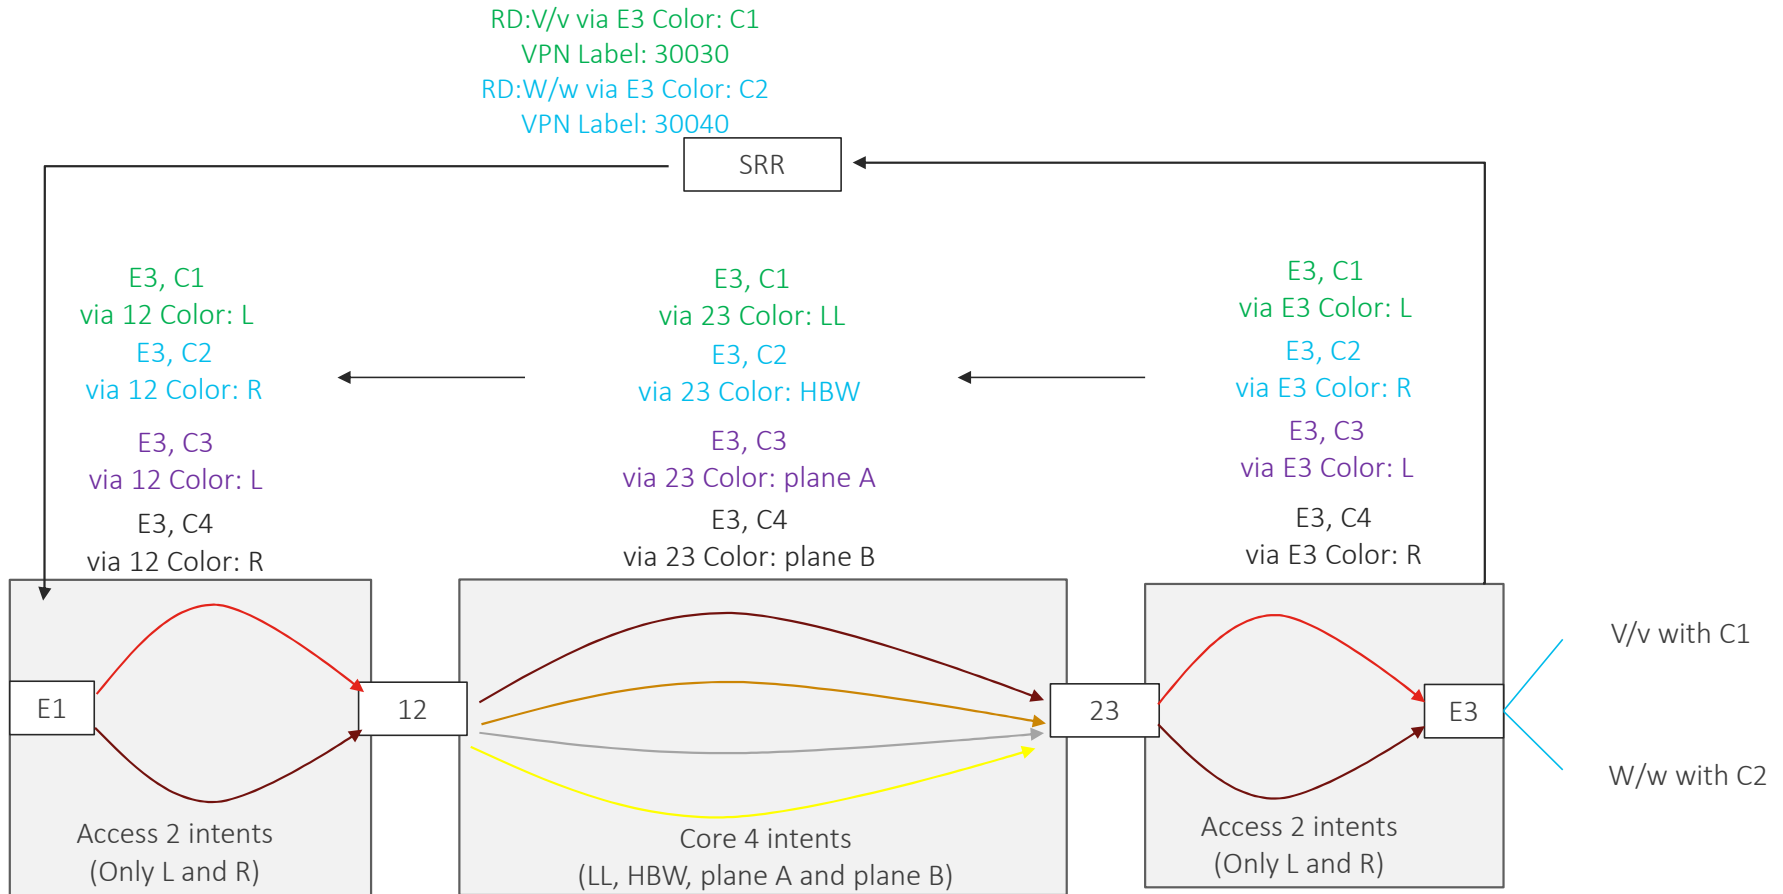
BGP-CAR: Inter-Domain Multiple color-aware paths

Base case: Intent-aware paths to a specific transport endpoint (e.g., PE loopback)
Many types of intent



BGP-CAR: N end to end intents over M intra domain intents (N>=M)

Case: Intent-aware paths across domains with different intent (color) granularity

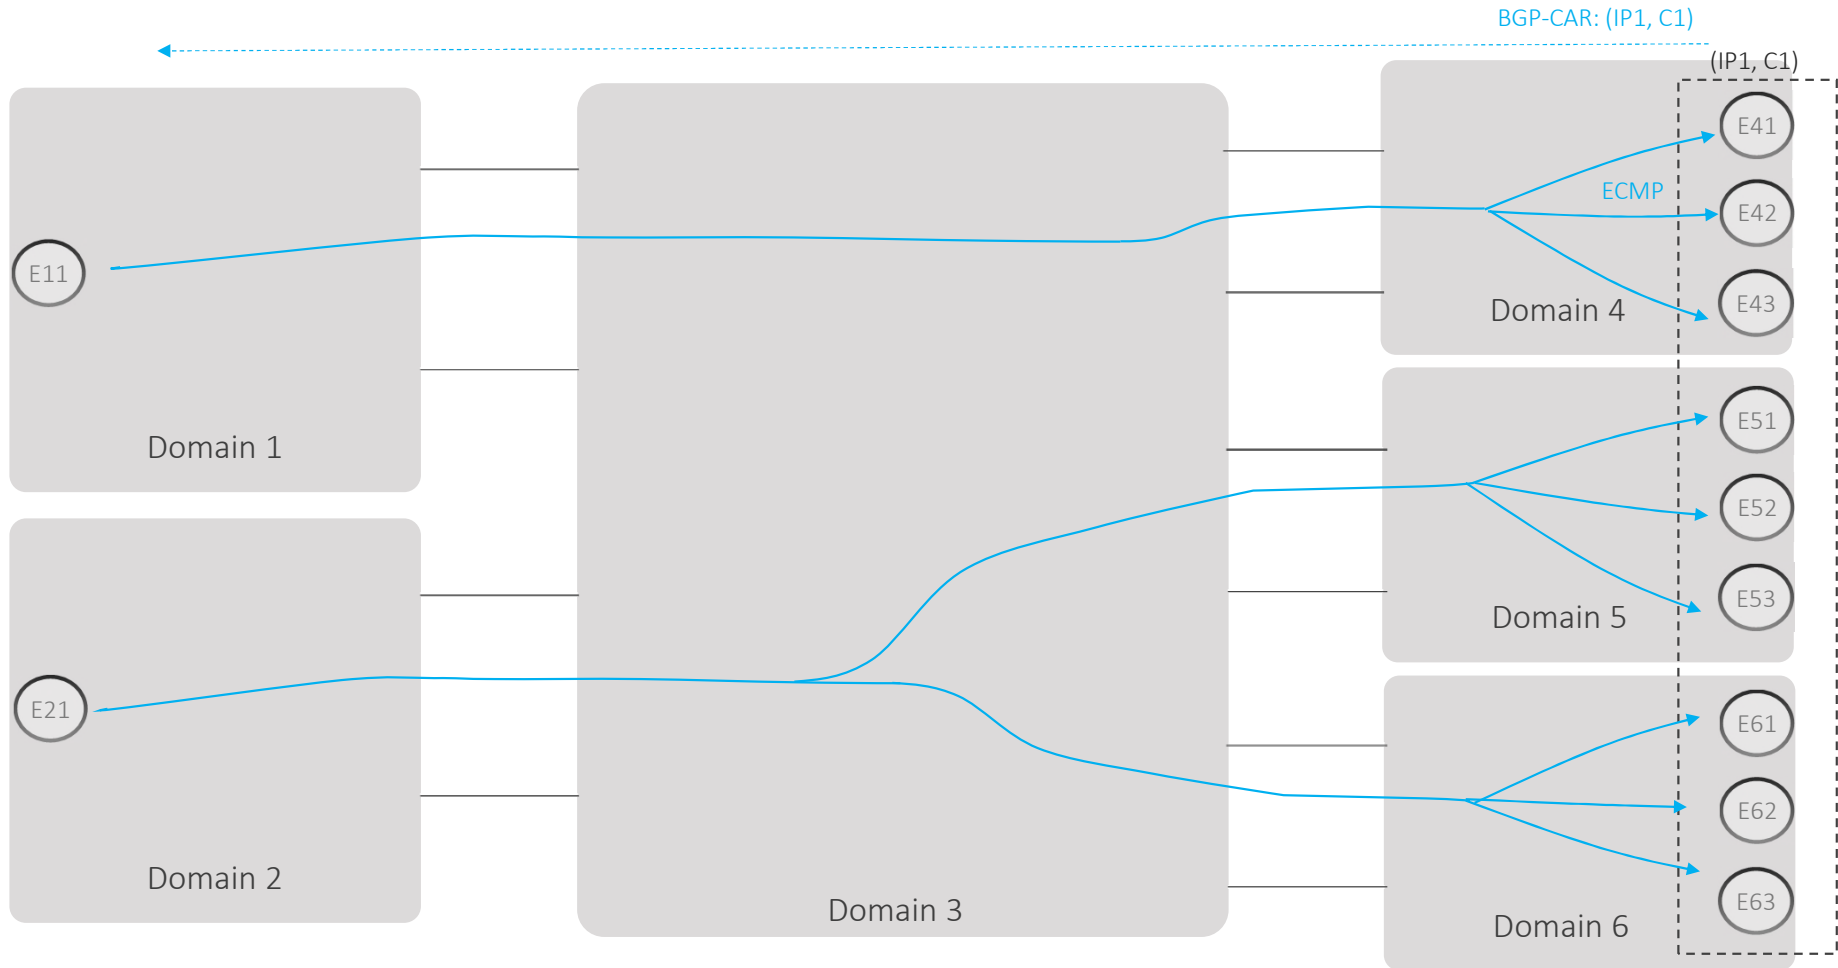


Service route steers over BGP CAR routes by attaching Color-EC. Similarly, BGP CAR routes may be resolved over a different intra-domain color by attaching Color-EC representing intra domain intent.

BGP-CAR: Transport Anycast

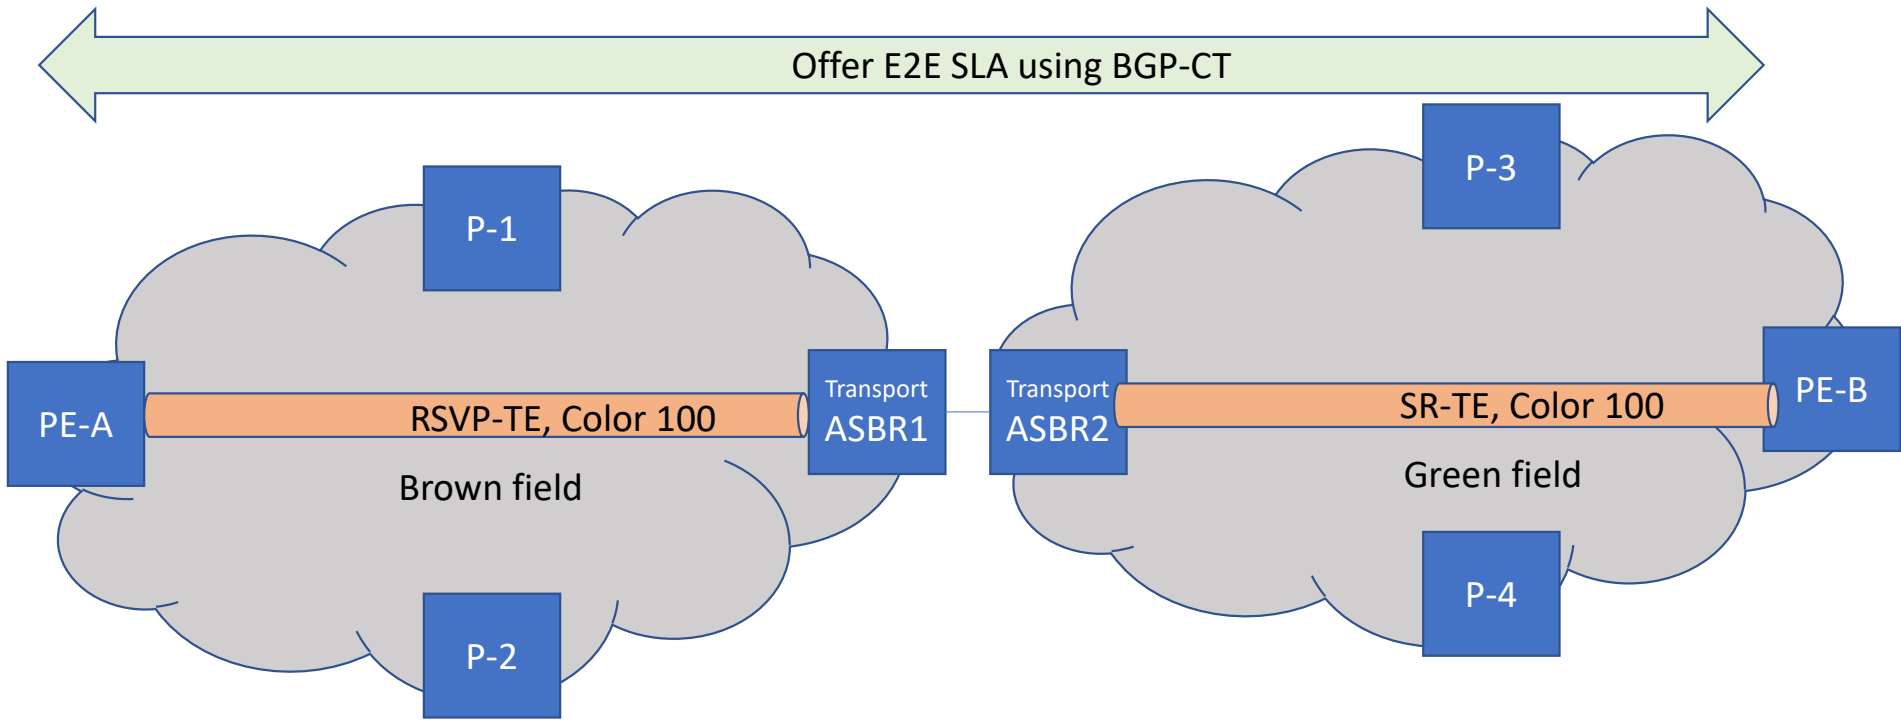
Case: Intent-aware paths to a transport Anycast IP

1. Network based Anycast
2. Ingress based Load-Balancing



BGP CT: Customer Use Case 1 – AT&T

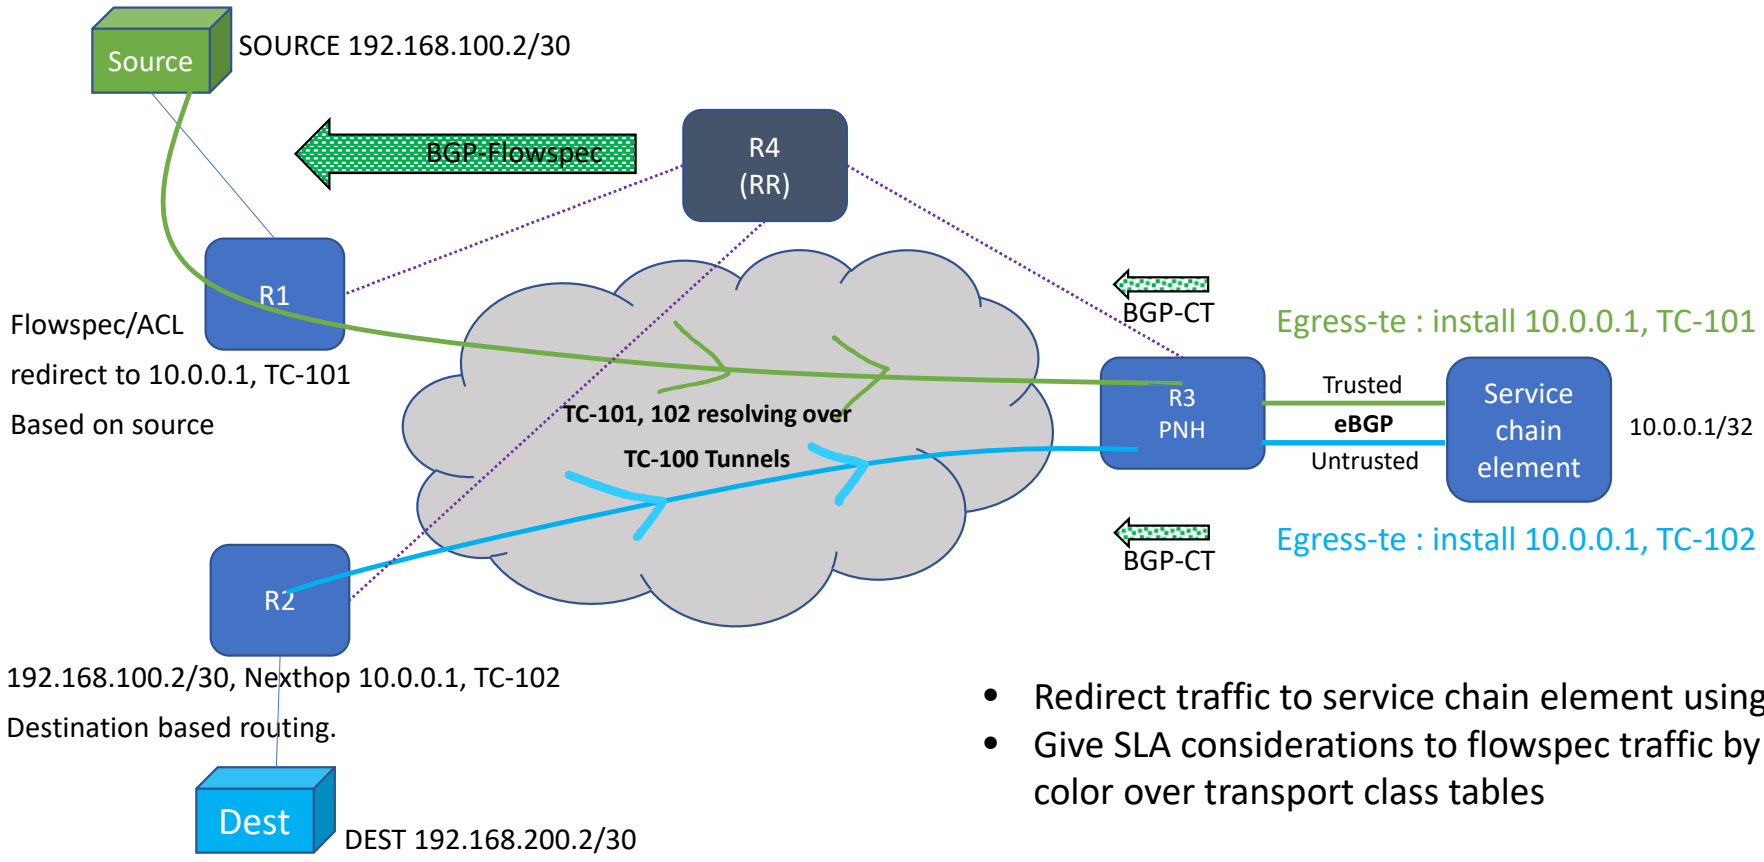
RSVP-TE/SR-TE coexistence during migration



- Current network is RSVP-TE
- Looking to introduce SR-TE in the newer network
- Need to maintain E2E SLA across both networks and RSVP-TE needs to understand color. Transport class provides route resolution accordingly

BGP CT: Customer Use Case 2

Offer SLA to Flowspec Traffic over RSVP-TE/SR tunnels



- Redirect traffic to service chain element using flowspec
- Give SLA considerations to flowspec traffic by resolving color over transport class tables

Flowspec redirect traffic to “BGP-CT EPE” end-points

BGP CT: Customer Use Case 3 – APAC

Network slicing across heterogenous color domains

“Currently, we have separate networks for domestic and international and they are independent so as the TE policies. We intentionally maintain the autonomy and modularity for administrative purposes. When we need inter-domain intent awareness, we would need the same level of flexibility in the proposed solution.

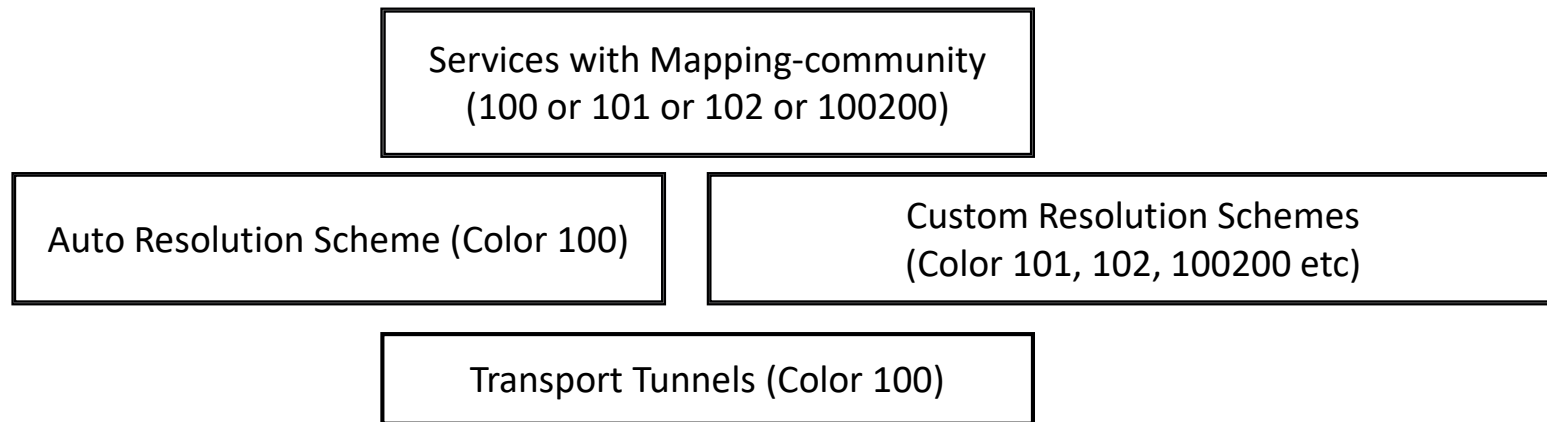
I would also like to highlight, service provider networks usually have more meshed paths in the core and aggregation domains where more granular intents can be realised. However, the access network domain will have less number of paths (either left or right in a ring / partial mesh / hub and spoke – in regional remote areas) where we would need only a few discrete transport classes / colours.

Hence, requirement for remapping of transport classes / colours within a single AS shouldn't be considered as a corner case in my opinion.”

- Moses Nagarajah (Telstra Networks)

BGP CT: Customer Use Case 3 – Solution

Heterogenous Color Granularity, Customizing resolution schemes



❑ Service Layer

- Carve out a service route mapping-community space across the AS domains
- Each mapping-community in this space is an “abstract value” identifying an SLA
 - (e.g “color:0:100200”, “color:0:400500”, “color:0:321332”)

❑ Transport Layer

- Customized resolution-schemes for BGP-CT family routes in relevant AS domains to use available colors
- Mapping-community **transport-target:0:101 (Medium Red)** and **transport-target:0:102 (Light Red)** can be
 - Custom Mapped to **tc-100 [Red]** in **AS Metro Domain A and C**
 - Strictly Mapped to **tc-101[Medium Red]** and **tc-102[Light Red]** respectively in **AS Core Domain B**

Issues raised in Adoption call Part 3

CAR

- Packing of PDUs
 - Compatible new NLRI (AFI/SAFI)
 - Key fields (color/NLRI)
 - Support for SRv6 + Robustness + operations
 - Scaling + Scalable replacement for LU
 - Intent at Service Level
 - Good BGP LU Follow-on
 - Use of Add Path (eBGP)
-
- CAR Anycast Scenario (A.7)
 - CAR color mapping resolution (effective)
 - CAR LCM and Color Communities
 - CAR Non-Agreeing color domains for Anycast End-Points
 - Filtering procedures/RTC (Sec 5)
 - ECMP, protection in Non-agreeing color domains (non anycast)

CT

- Packing of PDUs
 - Compatible new NLRI (AFI/SAFI)
 - Key fields (RD, NLRI, RT)
 - Support for SRv6 + Robustness + operations
 - Scaling + scalable replacement for LU
 - Intent at Service Level
 - Good BGP LU Follow-on
 - Use of Add Path (iBGP)
-
- CT Anycast Scenario Handling
 - SAFI-76 use only in Option C
 - Embedded MPLS label in 8277 addresses
 - Use of multiple RTs with non-agreeing Color domains
 - Multipath, local repair and churn suppression with RD in NLRI
 - Support for multiple encapsulations

CAR – First Things WG can help with

- Review and feedback for
 - Current version
 - Updates on SRv6 flows, considerations
 - Updates on filtering mechanisms
- Inputs on use-case scenarios

BGP-CT – First things WG can help with

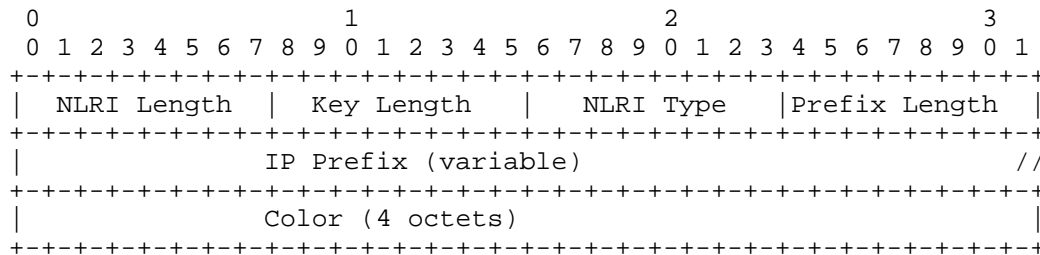
- ❑ Text clarification for disallowing SRv6 transposition for SAFI 76.
- ❑ Text clarification Section 8. usage for Route Distinguisher
 - Configuration flexibility for Same RD and Unique RD
 - Label allocation modes for BGP-CT
- ❑ Text for expressing and processing end customer intent (on CE-PE Links)
 - Control Plane Procedures for signaling customer Intent through provider networks end-to-end
 - Data Plane Procedures for mapping customer intent to provider network SLA
- ❑ Other Clarification of draft text (as required)

CAR-CT Interoperability Document

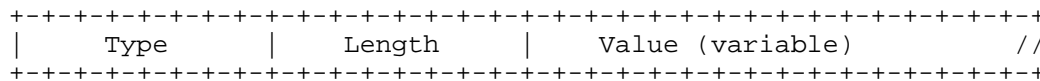
- Determine a path to interoperable features
 - Larger discussion, “what’s a color for?”
 - Color in the NLRI (CAR) and not in NLRI (CT)
 - Impact of Packing and translation of Packing
 - Error handling
 - Filtering – RT, RD, and use of RTC
 - Expressing transport intent/color for inter-as VPN scenarios

-CAR – NLRI and LCM community

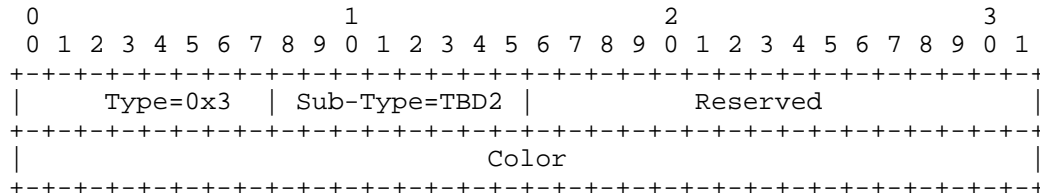
NLRI:



Followed by optional TLVs encoded as below:



LCM Extended Community:



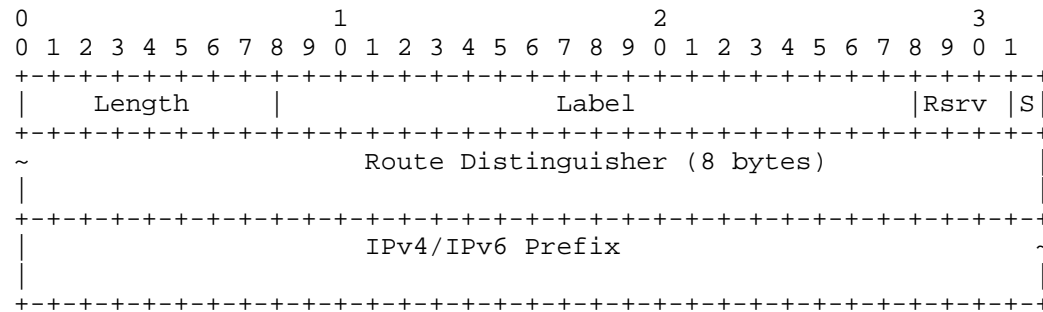
Defined Non-Key fields:

1. Label TLV (compare vs. BGP-LU)
2. Label Index TLV (alters BGP Prefix-SID behaviors)
3. SRv6 SID TLV (requires BGP Prefix-SID path attribute)

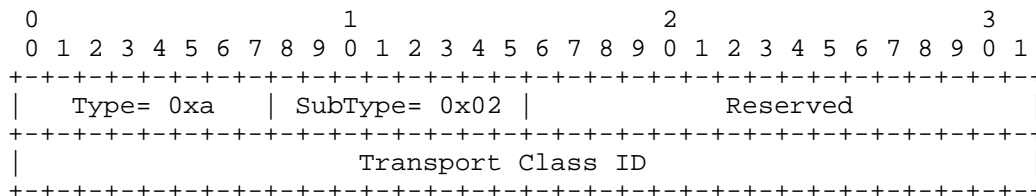
Local-Color-Mapping Extended Community added for crossing between color domains; overrides Color in NLRI for route resolution with color procedures. Optional (discouraged?) in other scenarios.

-CT – NLRI and transport community

NLRI:



Transport Class Extended Community:



Label (stack) compared to BGP-LU.

Label Index and SRv6 scenarios covered as per existing procedures applicable to BGP-LU using Prefix-SID.

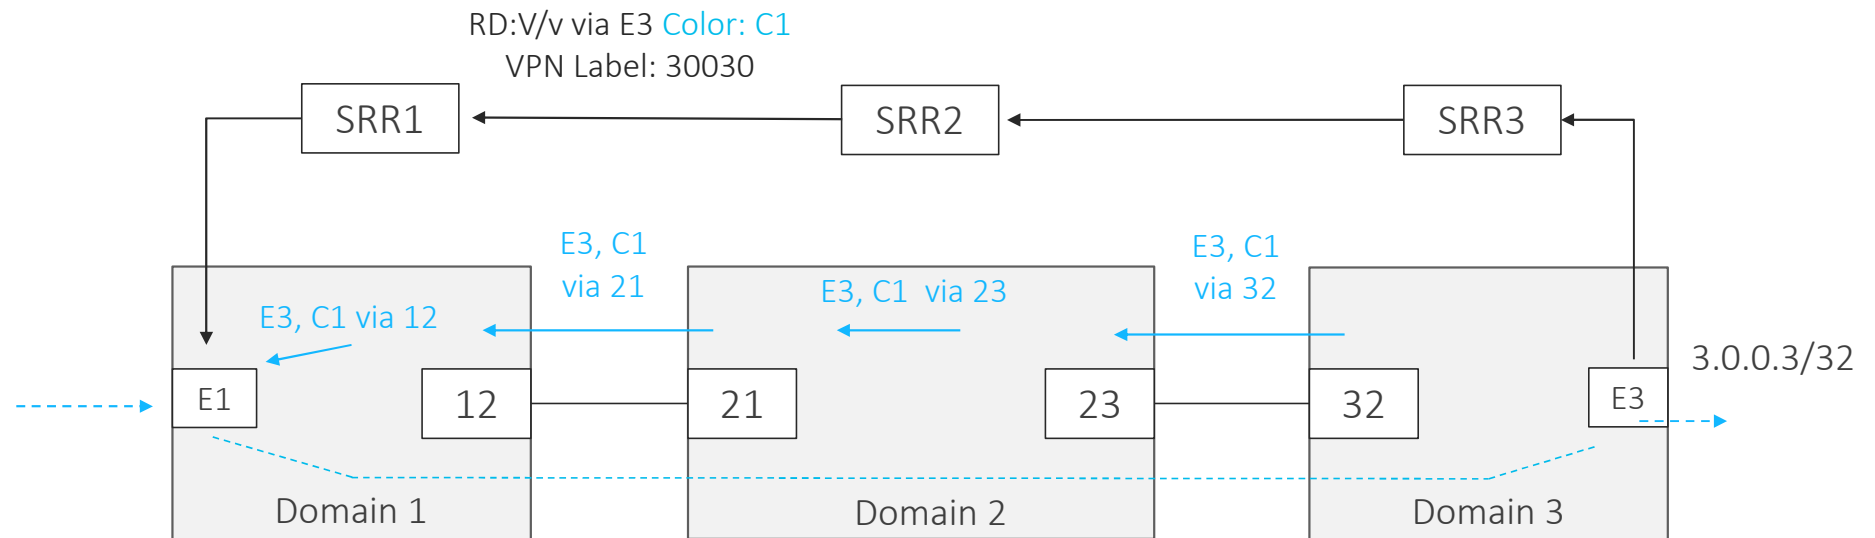
Backup slides for discussion

CAR backup slides

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
 - Color is the standard way to represent intent
 - Carried in BGP Color Extended-Community in BGP service routes (L3VPN, EVPN etc)

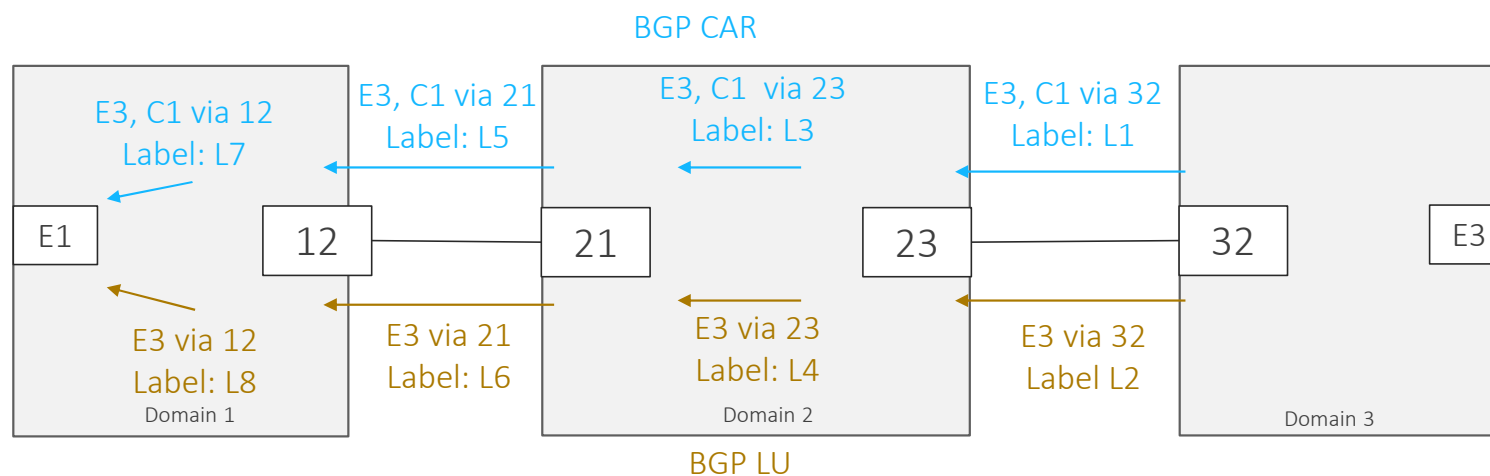
BGP Color-Aware Route & Automated Steering



- E3, C1 is a Color-Aware BGP route in underlay that provides intent-aware path to E3
- A C1 Colored service route RD:V/v from E3 is automatically steered onto a Color-Aware path (E3, C1)
 - Color: C1 carried in BGP Color Extended-Community attached to RD:V/v
- Steering for all services (L3VPN, EVPN, Internet/global table etc)

New SAFI in BGP

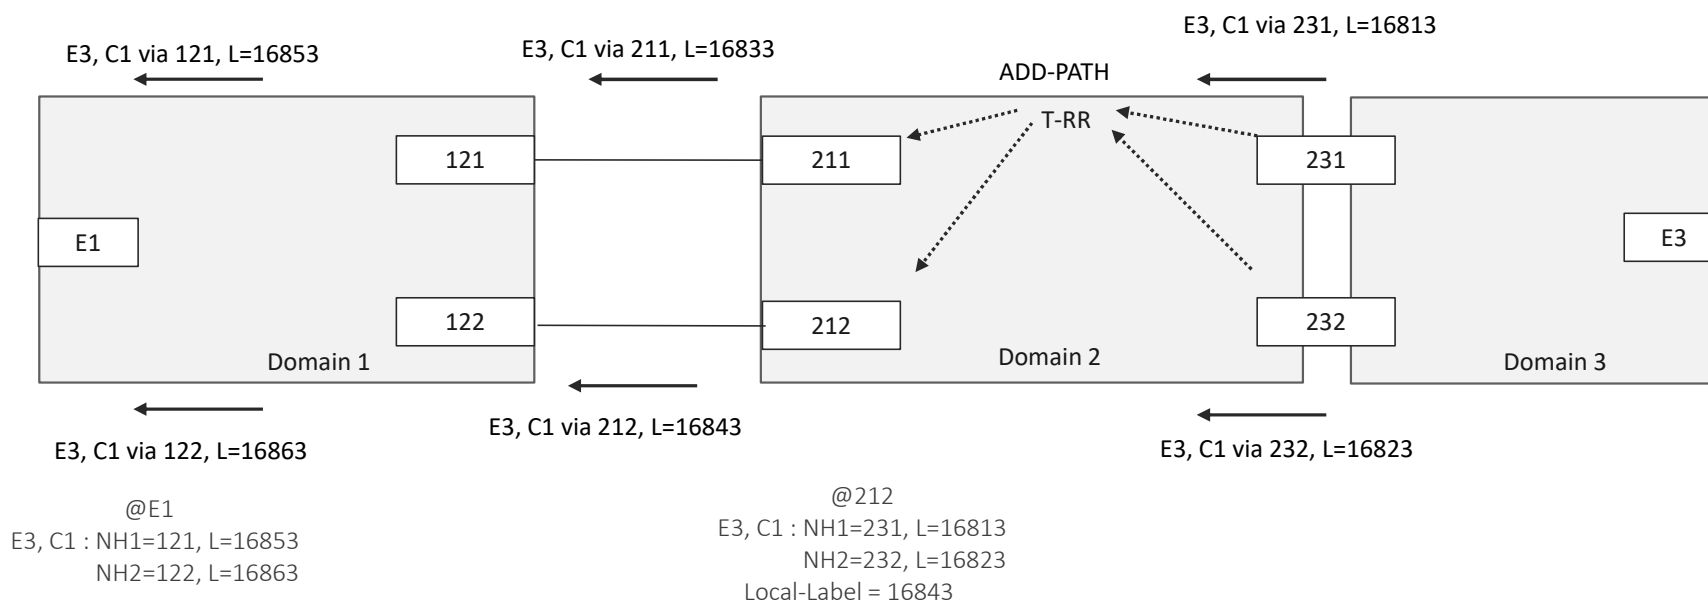
- Need ability to signal multiple instances of the same prefix for each color (i.e., intent)
- Evolution of best effort BGP-LU SAFI (RFC 3107/8277)
 - > Modernize, Address some of the limitations with BGP-LU
- Maintain functional and operational consistency with BGP-LU
- No need to use VPN constructs and machinery at every transport hop



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 each underlay hop
- Inherently provides ECMP-aware/backup paths at every hop
 - Faster, localized convergence
 - No need for VPN import to bring diverse path together with complex workarounds
- Most efficient for subscription
 - [E, C] direct lookup
- Consistent with SR Policy data model

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
- Note: BGP-CT fails to provide domain-local convergence and BR failure churn suppression

Extensible, Future-Proof NLRI Encoding

- New SAFI allows opportunity for better NLRI design
 - Existing SAFIs carry key (prefix) and non key information(eg: label in VPN,BGP-LU, EVPN)
 - > Hard coded in per SAFI specification
 - BGP CAR provides structure to this non-key information for future extensibility and flexibility
 - > No good reason to inherit constraints of current SAFIs, e.g., only a MPLS label field in NLRI
- Encode a NLRI (Route) Type
- Encode a key length
- Encode non-key TLVs
- Per route unique data in NLRI non-key TLVs; rest in Attribute
 - Provides packing efficiency for BGP updates

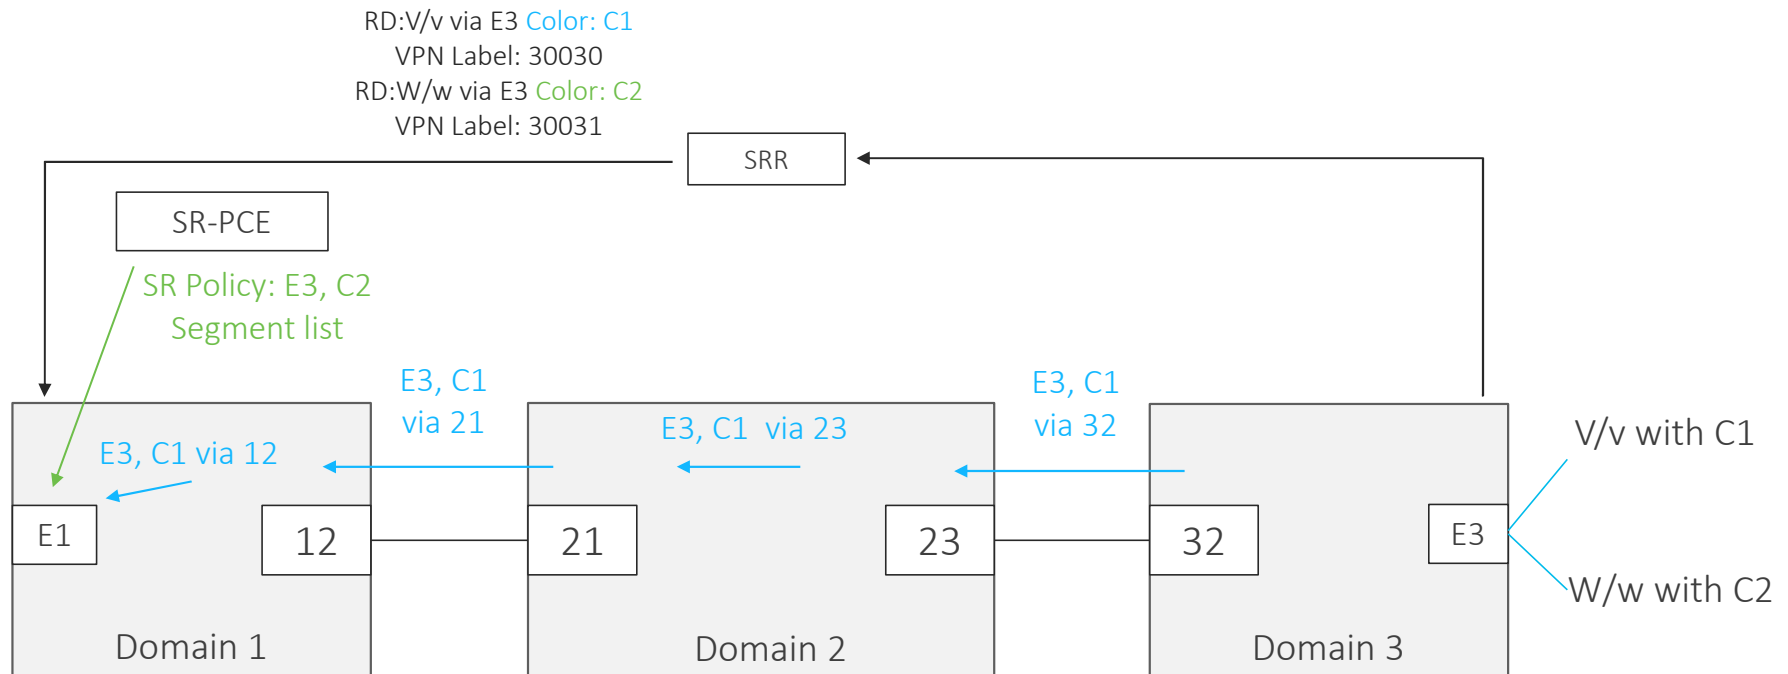
Encapsulations

- Ability to signal multiple encapsulations for a CAR route
 - Signaled via Non-Key TLVs
 - > MPLS Label(s), Label-Index, SRv6 SID(s) etc
 - Can signal separate label (or equivalent) values for different encapsulations
 - Efficient, preserves packing – e.g., label-index
 - Beneficial for co-existence, migration & interworking
 - > Efficient signaling, automated migration handling, operational simplicity
 - Avoids duplicate routes for each encap
 - Avoids separate control plane planes for distribution

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 will also be mapped to traditional mechanisms
 - RSVP-TE
 - IGP/LDP
 - BGP-LU
 - Supports brownfield, incremental deployment

Seamless BGP CAR and SR Policy co-existence with E,C model

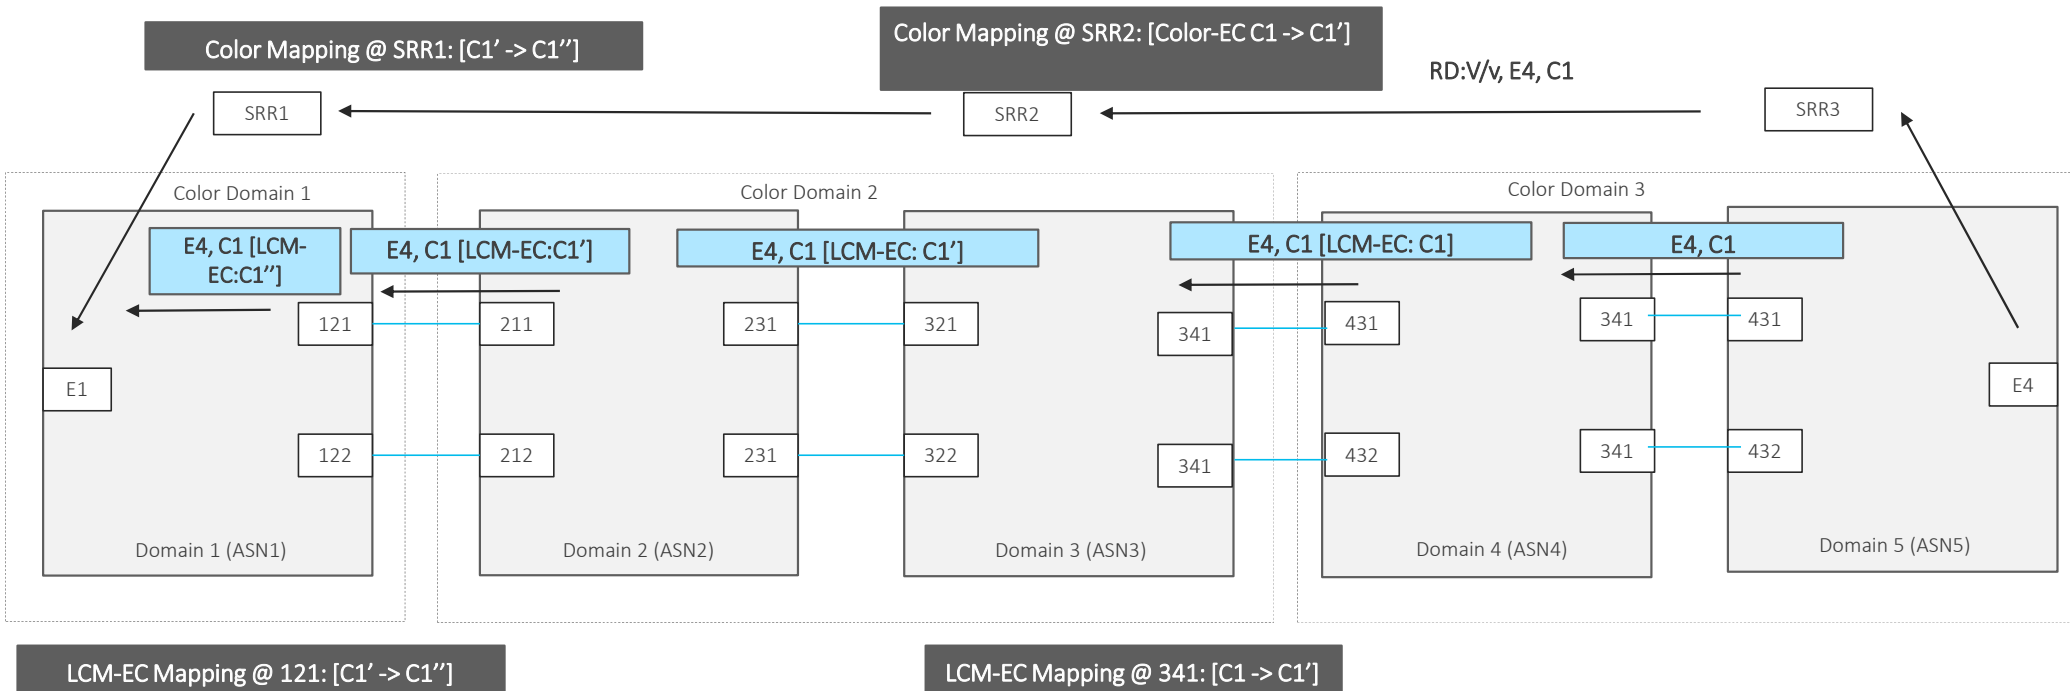


- E1 automatically steers the received colored service routes as follows:
 - V/v via (E3, C1) provided by BGP CAR
 - W/w via (E3, C2) provided by SR Policy
- Seamless compatibility with SR-Policy architecture - color, steering, fallback etc
 - Supported in multiple implementations, deployed

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 immutable, preserved e2e
 - Eases tracking of route
- E (Prefix) is unique in inter-domain transport network (e.g., PE)
 - Makes (E, C) unique e2e even if C is local to a color domain

Multiple Color Domains - Color Re-mapping



Additional aspects in draft

- CAR provides flexibility to operator for various dataplane designs, flat and hierarchical for scaling
 - CAR draft provides analysis of trade-offs and optimizations
- CAR design enables extension of intent-aware routing to PE-CE networks (VPN CAR)

Summary

- CAR is evolution of BGP-LU, but intent-aware
- CAR is extensible, built to accommodate new use-cases, multiple encapsulations efficiently
 - Defines a base framework that can be extended with low overhead
- CAR definitions focus on better protocol performance & scaling
 - Preserves update packing efficiency, memory storage
 - Avoids route duplication during migration
- CAR works seamlessly across traditional networks (LDP/RSVP-TE/BGP-LU)
- It is also totally compatible and consistent with SR-Policy/IGP-FA solution

IETF-114

BGP-CT – Backup Slides

Kaliraj, Nats

BGP-CT Recap. – Solution Constructs for Intra-Domain

- ❑ **Transport-class**: Collects tunnels with same TE characteristics (gold, Bronze, etc.). Identifier: 32-bit Transport-class ID.

- ❑ **Transport protocols** are made “transport-class aware”. RSVP-TE, SRTE, Flex Algo, etc.
 - ❑ Transport tunnels can be provisioned with a transport-class to install its ingress-routes in “Transport Route Databases” (RIBs used for resolution purposes only and not downloaded to the FIB) pertaining to the transport-class.

- ❑ **Resolution-scheme** : Realizes intended nexthop resolution to primary and fallback transport-class(es). Specifies an ordered list of primary and fallback Transport Route Databases. Identifier: Mapping-community.
 - ❑ Auto-created resolution-schemes for a transport-class, fallback on best-effort tunnels by default.
 - ❑ User-configured resolution-schemes fallback on other Transport-classes.

- ❑ **Mapping-community** : Service routes carry this community (e.g. Color:0:100) to signal their interest in a Resolution-scheme. Any BGP community/ext-comm can play this role.
 - ❑ Resolution of service-routes works the same way for Intra-AS and Inter-AS use cases.
 - ❑ The **first community on the route with a matching resolution-scheme** becomes the “**effective mapping-community**” for the route’s nexthop.

BGP-CT Recap. – Solution Constructs for Inter-Domain

- *How BGP disseminates transport-class information, across domain boundaries (ABR, ASBR).*

- ❑ **BGP-CT family:** A new transport-layer BGP family (SAFI 76),
 - ❑ advertises intra-AS transport tunnels to other domains, with indication of Transport-class.
 - ❑ follows RFC-4364 procedures, and RFC 8277 NLRI encoding.
- ❑ **Route Distinguisher:** Used to carry multiple tunnels to same destination, without BGP path-hiding at ASBRs and RRs.
- ❑ **Transport Route-Target:** route-target that identifies the transport-class of a BGP-CT route.
 - ❑ Leaks the route to the Transport Route Database of the transport-class after stripping the RD.
 - ❑ “transport-target:0:<n>” acts as a mapping-community that resolves a BGP-CT route strictly over transport-class <n>, with no fallback(s). This preserves the SLA <n> end-to-end in the transport-layer.
 - ❑ A new IANA RT format type, to avoid collisions with service-plane route-targets. To ease deployment.
- ❑ Transport Route Databases are populated with BGP-CT routes, for inter-AS service-endpoints.
- ❑ Option-C style Label-swap for BGP-CT routes at ABR and ASBR preserves the transport-class end-to-end in the inter-AS network.

Customer Use Case – Network Slicing

“Currently, we have separate networks for domestic and international and they are independent so as the TE policies. We intentionally maintain the autonomy and modularity for administrative purposes. When we need inter-domain intent awareness, we would need the same level of flexibility in the proposed solution.

I would also like to highlight, service provider networks usually have more meshed paths in the core and aggregation domains where more granular intents can be realised. However, the access network domain will have less number of paths (either left or right in a ring / partial mesh / hub and spoke – in regional remote areas) where we would need only a few discrete transport classes / colours.

Hence, requirement for remapping of transport classes / colours within a single AS shouldn't be considered as a corner case in my opinion.”

- Moses Nagarajah (Telstra Networks)

https://mailarchive.ietf.org/arch/msg/idr/x9Zy9ob5_78bsAiE5Pvr9qAJIKw/

BGP-CT – Solution

Heterogenous Color Granularity

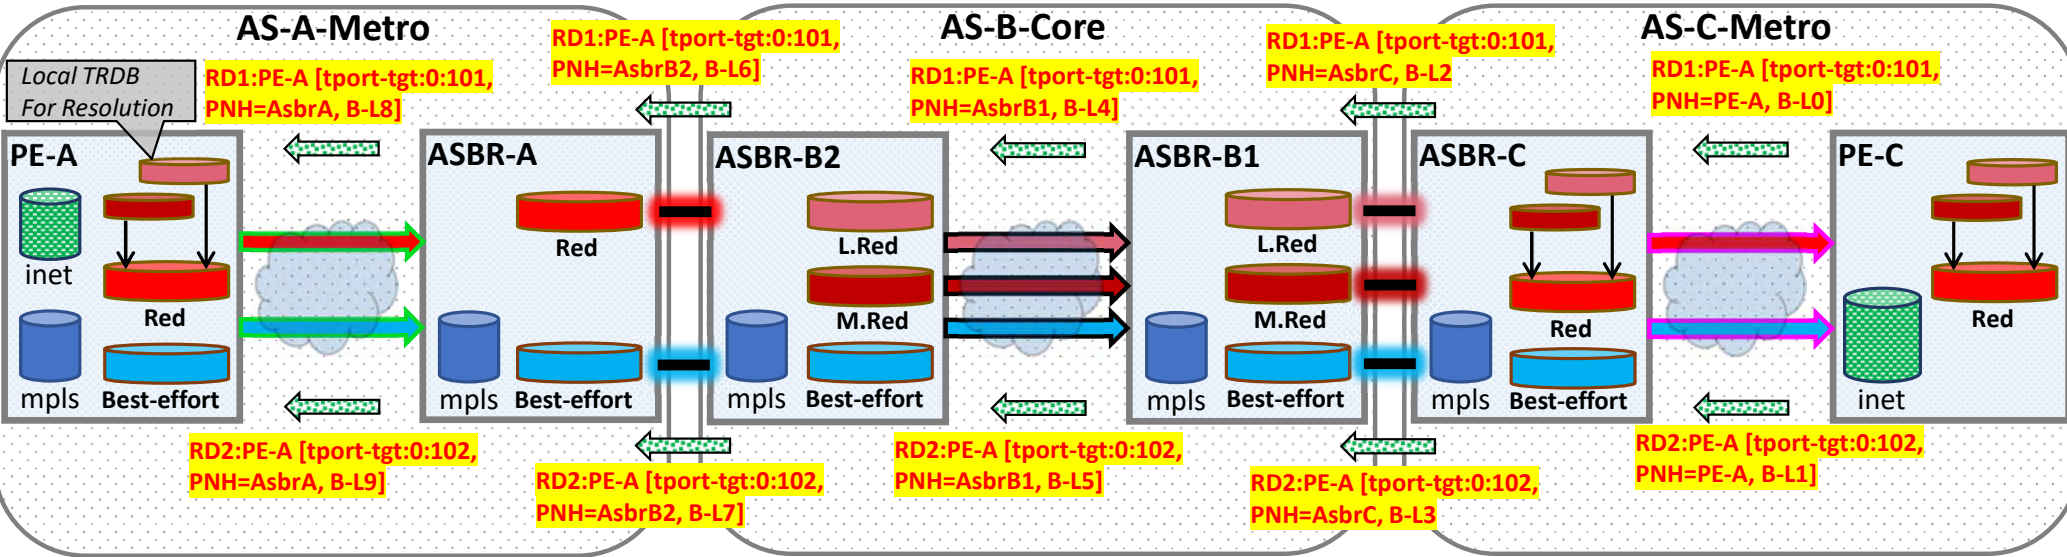
☐ Service Layer

- Carve out a service route mapping-community space across the AS domains under the same admin control
- Each mapping-community in this space is represented by an “abstract value” and identifies an SLA
 - (e.g “color:0:100200”, “color:0:400500”, “color:0:321332”)
- Service routes attach mapping-communities **ONLY from this space**

☐ Transport Layer

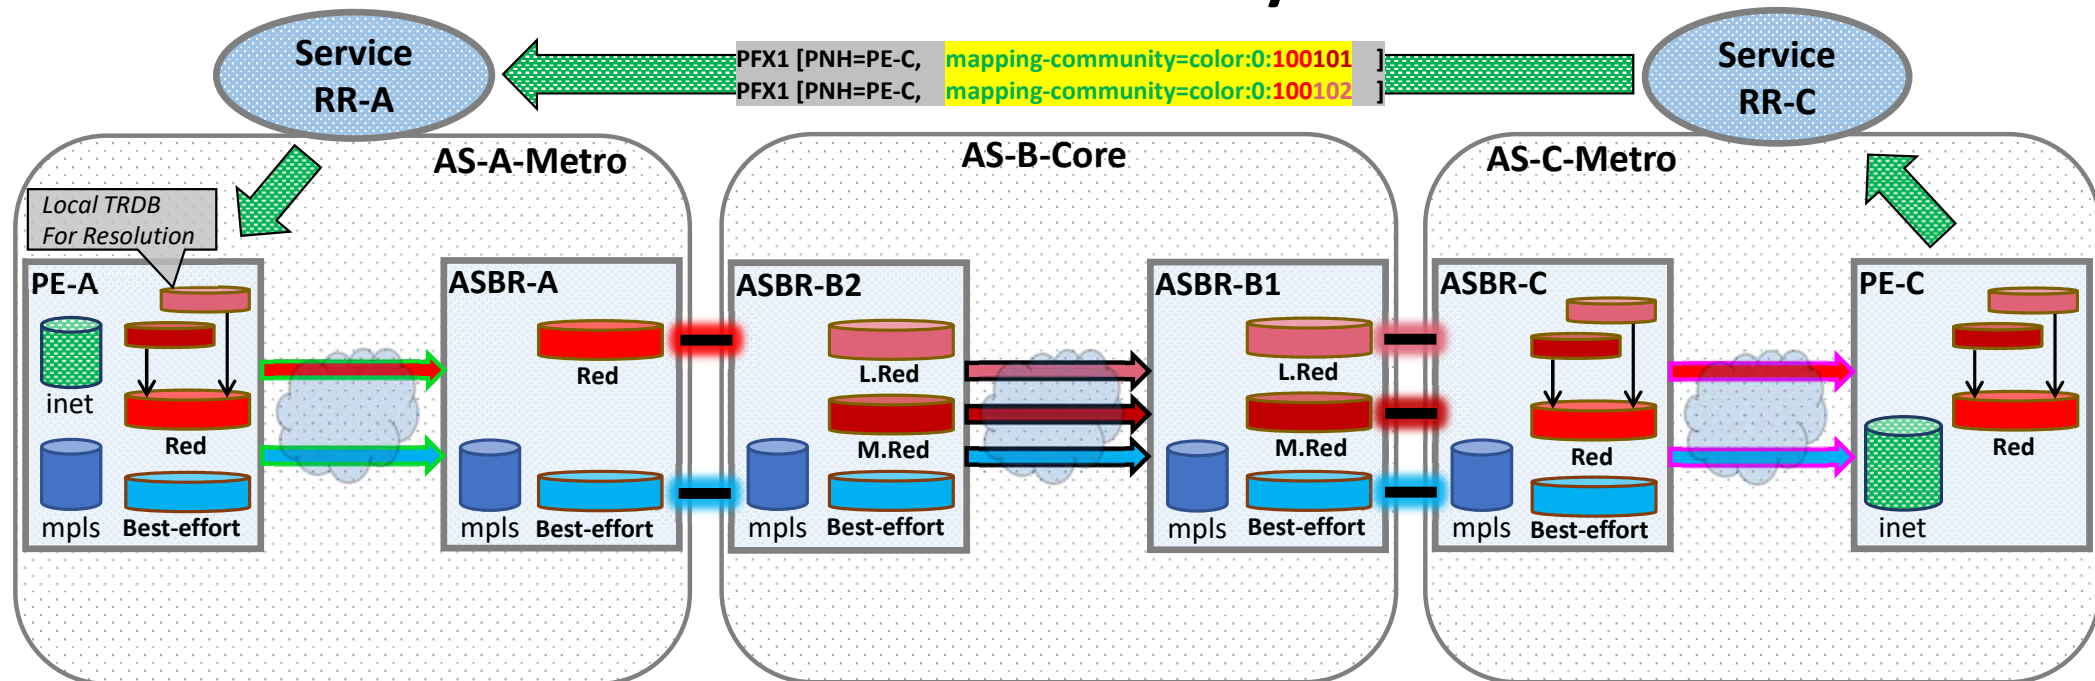
- Customized resolution-schemes for BGP-CT family routes in relevant AS domains
- Model the BGP CT transport route-target as shades of the same color
 - (e.g 100:Red, 101:Medium-Red, 102:Light-Red)
- Customize resolution-schemes to map the same transport route-target mapping-community to available shades
- Mapping-community **transport-target:0:101 (Medium Red)** and **transport-target:0:102 (Light Red)** can be
 - Custom Mapped to **[Red]** in **AS Metro Domain A and C**
 - Strictly Mapped to **[Medium Red]** and **[Light Red]** respectively in **AS Core Domain B**
- **Pros:** No rewrites required at domain boundaries, Service Prefix level unique shading end-to-end
- **Cons:** More resolution schemes in AS Metro Domains

BGP-CT – Solution Illustration – Transport Layer



PE-A Transport Classes		ASBR-B2 Transport Classes		ASBR-C Transport Classes		LEGEND			
<input type="checkbox"/> Red: 100		<input type="checkbox"/> M.Red: 101 <input type="checkbox"/> L.Red: 102		<input type="checkbox"/> Red: 100			— Inter-AS Link → RSVP-TE ← SAFI 76		
PE-A Transport Resolution Schemes		ASBR-B2 Transport Resolution Schemes		ASBR-C Transport Resolution Schemes				→ SR-TE → FLEX ALGO	
M.Red_Custom L.Red_Custom <input type="checkbox"/> map-comm trt:0:101 <input type="checkbox"/> map-comm trt:0:102 <input type="checkbox"/> [Red] <input type="checkbox"/> [Red]		M.Red_Strict L.Red_Strict <input type="checkbox"/> map-comm trt:0:101 <input type="checkbox"/> map-comm trt:0:102 <input type="checkbox"/> [M.Red] <input type="checkbox"/> [L.Red]		M.Red_Custom L.Red_Custom <input type="checkbox"/> map-comm trt:0:101 <input type="checkbox"/> map-comm trt:0:102 <input type="checkbox"/> [Red] <input type="checkbox"/> [Red]					
PE-A M.Red TRDB		ASBR-A mpls FIB		ASBR-B2 mpls FIB		ASBR-B1 mpls FIB		ASBR-C mpls FIB	
PFX1 [PUSH B-L8 PUSH ASBR-A-red]		B-L8 [SWAP B-L6 PUSH ASBR-A-red-intf]		B-L6 [SWAP B-L4 PUSH ASBR-B1-M.Red]		B-L4 [SWAP B-L2 PUSH ASBR-C-M.red-intf]		B-L2 [SWAP B-L0 PUSH PE-C-Red]	
PE-A L.Red TRDB									
PFX2 [PUSH B-L9 PUSH ASBR-A-red]		B-L9 [SWAP B-L7 PUSH ASBR-A-red-intf]		B-L7 [SWAP B-L5 PUSH ASBR-B1-L.Red]		B-L5 [SWAP B-L3 PUSH ASBR-C-L.red-intf]		B-L3 [SWAP B-L1 PUSH PE-C-Red]	

BGP-CT – Solution Illustration – Service Layer



PE-A Service Resolution Schemes

Srvc_M.Red_BE

□ map-comm color:0:100101

□ [L.Red, BestEffort]

Srvc_L.Red_BE

□ map-comm color:0:100102

□ [M.Red, BestEffort]

PE-A inet FIB

PFX1 [PUSH B-L8
PUSH ASBR-A-red]

PFX2 [PUSH B-L9
PUSH ASBR-A-red]

ASBR-A mpls FIB

B-L8 [SWAP B-L6
PUSH ASBR-A-red-intf]

B-L9 [SWAP B-L7
PUSH ASBR-A-red-intf]

ASBR-B2 mpls FIB

B-L6 [SWAP B-L4
PUSH ASBR-B1-M.Red]

B-L7 [SWAP B-L5
PUSH ASBR-B1-L.Red]

ASBR-B1 mpls FIB

B-L4 [SWAP B-L2
PUSH ASBR-C-M.red-intf]

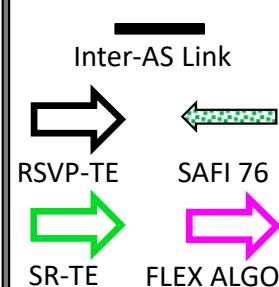
B-L5 [SWAP B-L3
PUSH ASBR-C-L.red-intf]

ASBR-C mpls FIB

B-L2 [SWAP B-L0
PUSH PE-C-Red]

B-L3 [SWAP B-L1
PUSH PE-C-Red]

LEGEND



BGP-CT

Action Items for the working group

- 2nd Implementation for BGP-CT
- Text clarification for disallowing SRv6 transposition for SAFI 76.
- Text clarification Section 8. usage for Route Distinguisher
 - Configuration flexibility for Same RD and Unique RD
 - Label allocation modes for BGP-CT
- Text for expressing and processing end customer intent (on CE-PE Links)
 - Control Plane Procedures for signaling customer Intent through provider networks end-to-end
 - Data Plane Procedures for mapping customer intent to provider network SLA
- Other Clarification of draft text (as required)

CAR-CT Comparability

Jeff Haas

Comparability Points

- Color in the NLRI or not?
 - Mostly caters to each proposal's default deployment models.
 - Both procedures are required to deal with color in an Extended Community. (See regular commentary, "VPN Import Procedures".)
 - RD isn't color. RD could carry color as an option?
- Larger discussion, "what's a color for?"
 - Point of assignment(s) and use varies for default scenarios in each proposal.
 - Choices impact level of redundancy of BGP routes in network or need for deploying add-paths selectively for providing access to redundant paths.
 - Each proposal can achieve similar forwarding as the other proposal by making intentional choices in color/RD assignments.

Comparability Points

Packing impacts:

- Both proposals pack label stacks similarly. -CAR imposes new transitivity requirement on label stack in its encoding.
 - Transitivity flag helpful for situations where Prefix-SID must be present!
- Label Index denser in -CAR.
- In RFC 9252 encoding, for -CAR entire SID could be encoded (1..16 bytes). -CT follows existing SRv6 procedures with labels.

Error Handling:

- -CAR optional NLRI component introduces potential RFC 7606 considerations. RFC 7606 adjusted BGP protocol procedures to ensure NLRI was more easily parseable vs. Path Attributes to permit Treat-As-Withdraw procedures.

Comparability Points

- Filtering:
 - -CT uses RT-Constrain vs. Transport Class extended community. Thus, filter is vs. routes containing a particular color.
 - -CAR's procedures are not fully specified. Desired filter behavior is (E,C). Is this C effective color or color in NLRI?
 - Both procedures may have multi-color domain filtering considerations since effective color may change at a domain boundary.

Comparability Points

- Expressing transport intent/color for inter-as VPN scenarios:
 - -CAR has VPN-CAR. Basically CAR NLRI + RD with route targets.
 - -CT has no documented procedures for this scenario.

Comparability Points

- Expressing transport intent/color for inter-as VPN scenarios:
 - -CAR has VPN-CAR. Basically CAR NLRI + RD with route targets.
 - -CT has no documented procedures for this scenario.