

Shepherd Review of CAR WG LC Issues

- CAR-WGLC-RTG-DIR-Review-1 3
 - Original Text: 3
 - Email Text : 3
 - Discussion: 3
 - Shepherd’s Action item: 3
- CAR WGLC Q1 – Issue 1 - Motivation for IP Prefix 4
 - Original Text: 4
 - Shepherd’s Review of CAR WGLC Q1 – Issue 1 4
 - Issue 1.1: What NLRIs does drat-ietf-idr-bgp-car-02 define in CAR SAFI and VPN CAR SAFI? 5
 - Issue 1.2: Are the NLRIs defined by this document for the CAR SAFI and CAR VPN SAFI passing only infrastructure routes? 5
 - Issue 1.3: What SRv6 SIDs get passed in CAR SAFI and VPN SAFI? 6
 - Issue 1.4 - Why should CAR SAFI and VPN CAR SAFI have type 2 routes? 7
 - Issue 1.5 – Proof based on section 9.1.1? 9
 - Shepherd Concern: Redistribution of Customer Routers 10
 - Issue 1 - Action item summary for issues: 10
- Shepherd’s Editorial Comments Related to Issue-1 11
 - Editorial Comment 1: 11
 - Editorial Comment 2: 11
 - Editorial Comment 3: Section 2.1 11
 - Editorial Comment 4, Section 9.1.2 11
 - Editorial Comment 4, Section 9.1.2 11
- CAR-WGLC-Q1-Issue-2: Why Add IP Prefix? 12
 - Original Text: 12
 - Shepherd’s action items: 13
- CAR-WGLC-Q1-Issue-3 – Clarity of Support for IP Prefix in AFI/SAFI Section 10 + No Implementations ... 14
 - Original Text: 14
 - Shepherd’s comments 16
 - Shepherd’s action items: 16
- CAR-WGLC-Q1-Issue-5 – Lack of Clarity for Procedures in Sections 8-10 17
 - Original Text: 17
 - Shepherd’s comments: 18

Shepherd's Report for CAR WG LC Issues

Shepherd's action items:.....	18
CAR-WGLC-Q1-Issue-6: Relationship among multiple drafts related to SRV6	19
Original text:.....	19
Action item:.....	19
CAR-WGLC-Q1-Issue-7 – Using CAR SAFI and CAR VPN SAFI for transport	20
Original Text:	20
Shepherd's action items:.....	20
CAR-WGLC-Issue-8 - Use of CAR-SAFIs for Services Families	21
Original text:.....	21
Shepherd's Action items:	21
OPS-DIR review	22

CAR-WGLC-RTG-DIR-Review-1

Original Text:

Link: https://mailarchive.ietf.org/arch/msg/idr/OX_q_e09ejtb0NZQ25smUrklwmQ/

Reviewer: Ben Niven-Jenkins

Review result: Has Issues

Email Text :

I have been selected to do a routing directorate "early" review of this draft.
<https://datatracker.ietf.org/doc/draft-ietf-mpls-rfc3107bis/>

The routing directorate will, on request from the working group chair, perform an "early" review of a draft before it is submitted for publication to the IESG. The early review can be performed at any time during the draft's lifetime as a working group document. The purpose of the early review depends on the stage that the document has reached.

For more information about the Routing Directorate, please see
<https://wiki.ietf.org/en/group/rtg/RtgDir>

Document: draft-ietf-idr-bgp-car-02

Reviewer: Ben Niven-Jenkins

Review Date: 31 July 2023

Intended Status: Experimental

Summary:

I have some minor concerns about this document that I think should be resolved before it is submitted to the IESG.

Comments:

The document provides lots of details in the main sections and examples in the appendices. I think that is generally ready to be published. The only suggestion I would have is to provide a short overview introducing BGP CAR and color aware routing before leaping into terminology and the details of the BGP CAR SAFI. Expanding the introduction section with an extra paragraph or two would be sufficient.

Discussion:

DJ-R-1: We appreciate your review of the CAR draft. We will add some overview text as you suggest.

<https://mailarchive.ietf.org/arch/msg/idr/ODL4N39sIPOqICJUMo7UITkdCUI/>

Shepherd's Action item:

1. Add Short overview of BGP CAR and Color-Aware routing in section 1.

CAR WGLC Q1 – Issue 1- Motivation for IP Prefix

Original Text:

- **Moshiko-R-1: (Response to DJ-R1) – Motivations for change**
 - So far CAR has advocated having Color in the NLRI. Here in the new prefix in Section 10, there is no Color in the NLRI. Which seems like a major change in ideology. Is there any problem with Color in the NLRI that you now want it [Color] in BGP community?
 - **Link:** <https://mailarchive.ietf.org/arch/msg/idr/o8I1dOTiLc-ImelfQUij1pOOJk/>
 - **DJ-R-2:** Hello Moshiko, Thank you for your query.
 - **Link:** https://mailarchive.ietf.org/arch/msg/idr/ayRYqImVVro48_AfqYSh4lukBR4/
 - There is no change in the existing specified semantics of the CAR (E, C) NLRI.
 - The IP prefix route type is an extension for the specific case where IP Prefix == intent or color, such as in SRv6 where a locator is assigned for a given intent. Here there is no requirement to have multiple instances of the prefix with different colors, hence no need to have the color in the NLRI.
 - As described in Section 9.1.1 for SRv6, the usage for the IP prefix route to resolve the Service SID via an underlay routed locator prefix or summary route (ref. Section 5 of RFC 9252). This is equivalent to the locator prefix being distributed in IGP Flex-Algo.
 - Being a regular BGP IPv6 route as specified, it follows the exact semantics as RFC 4271 and RFC 2545, allowing IPv6 routing/forwarding and summarization for these prefixes.
 - But by using a distinct CAR SAFI to distribute these transport routes, we avoid overloading of the IPv6 unicast SAFI (2/1) which also carries Internet service routes, as described in Section 9.3. Hope this clarifies.

Shepherd's Review of CAR WGLC Q1 – Issue 1

The abstract of draft-ietf-idr-bgp-car-00 states

This document describes a BGP based routing solution to establish end-to-end intent-aware paths across a multi-domain service provider transport network. This solution is called BGP Color-Aware Routing (BGP CAR). (*see Editorial Note 1)

This leads one to believe that enabling intent is the purpose of this draft. Section 2.1 describes the BGP CAR data model as encoding of the following in the NLRI:

- NLRI Key: IP Prefix, Color
- NLRI non-Key fields: MLS label stack, Label Index, SRv6 SID list, etc. [*See Editorial note 2]

Also included in the BGP update packet with issue is: BGP Next Hop, AIGP Metric, and Local-Color-Mapping Extended Community (LCM-EC). [See editorial note 3]

Section 2.2 on extensible encoding restricted to Color makes sense in that [E, C] becomes a new prefix type. However, it appears based on the email threat that the authors intend the following:

“The IP prefix route type is an extension for the specific case where IP Prefix == intent or color, such as in SRv6 where a locator is assigned for a given intent. Here there is no requirement to have multiple instances of the prefix with different colors, hence no need to have the color in the NLRI.” (DJ, https://mailarchive.ietf.org/arch/msg/idr/ayRYqImVVro48_AfqYSh4lukBR4/)

There are four points that need to be clarified by the CAR authors:

1. What NLRIs are included in draft-ietf-idr-bgp-car-02 definition CAR and CAR VPN services
2. Are the NLRIs defined by this document for the CAR SAFI and CAR VPN SAFI passing only infrastructure routes?
3. What SRv6 SIDs get passed in CAR SAFI and VPN SAFI?
4. Why should CAR SAFI and VPN CAR SAFI have type 2 routes?
5. Whether 9.1.1 provides enough proof for Type 2 NLRI.

Issue 1.1: What NLRIs does draft-ietf-idr-bgp-car-02 define in CAR SAFI and VPN CAR SAFI?

Text statements: draft-ietf-idr-bgp-car-02 specifies in section 2.9 that the CAR SAFI applies only to AFI 1 and 2. The CAR VPN SAFI is specified in section 8 with an example, and states:

“VPN CAR distribution for (RD, V, CC) requires a new SAFI that follows the same VPN semantics as defined in [RFC4364], the difference being that the advertised routes carry [the] CAR NLRI defined in section 2.9.2.”

This short discussion implies VPN CAR SAFI is valid for AFI=1 and SAFI=2, but the draft does not clearly state it in the specification.

Shepherd's action items to address Problem 1:

- The abstract should include the fact that two new SAFIs are being created.
- Section 1 paragraph 1 should indicate that this draft specifies two new SAFIs (CAR (83) and CAR VPN (84)). This paragraph should clearly state which AFIs these two SAFIs apply to. This should include AFI/SAFIs equal to 1/83, 2/83, 1/84, 2/84.
- Section 2.1 discussion of the database should indicate this document only defines the CAR SAFI for AFI=1 and AFI=2.
- Section 8 should clearly define which AFIs this document defines the VPN CAR SAFI.

Issue 1.2: Are the NLRIs defined by this document for the CAR SAFI and CAR VPN SAFI passing only infrastructure routes?

Shepherd's comment: All IDR drafts for intent-aware routing focuses on creating colored infrastructure paths (draft-ietf-idr-bgp-car, draft-ietf-idr-bgp-ct, draft-wang-idr-cpr). draft-ietf-idr-bgp-car-02 distinguishes between service routes and infrastructure routes. Section 3 of draft-ietf-idr-bgp-car-02 on service routes states the following:

“E1 automatically steers a C-Color aware service route V/v from E2 onto a (E2, C) color-aware path. If several such paths exist, a preference scheme is used to select the best path. This is consistent with the automatic service route steering on SR Policy (a routing solution providing color-aware path) defined in [RFC9256].

From this discussion, the shepherd concludes that the service routes are being passed in other NLRIs than the CAR NLRI (1/83, 2/83) or VPN CAR NLRI (1/84, 2/84). Based on the text, one could imply that BGP Policy steers the C-Color aware service route to (E, C) of CAR VPN.

The normative text does not identify which AFI/SAFIs can be supported as service routes by this document. The examples in Appendix A provide non-normative examples of IPv4 (AFI/SAFI 1/1) or IPv6 (AFI/SAFI) service routes. Section 3 lists the supported SAFIs by name, but these should be listed by name (AFI/SAFI) in section 3.

Section 8 of draft-ietf-idr-bgp-car-02 states the following:

“VPN CAR distribution for (RD, V, C) requires a new SAFI that follows [the] same VPN semantics as defined in [RFC4364], the difference being that the advertised routes carry CAR NLRI defined in section 2.9.2 of this document.”

Therefore, the infrastructure routes with color (E, C) are being passed in a CAR VPN (RD, V, C) to distinguish between two VPN customers using the same (E, C) routes.

It appears from Appendix A that service routes are being passed in IPv4 NLRIs (AFI=1/SAFI=2) and AFI=2/SAFI=2).

Shepherd questions:

- What NLRIs are being used to pass service routes?
- What procedures tie the service routes to the infrastructure routes?

Shepherd actions:

- Sections 3, 8, 9 and 10 – need to provide a consistent picture of what NLRIs are supported for service routes and infrastructure routes for intent-based/color aware routes.

Shepherd's comment:

Section sections 8-10 position in the document lead the reader to suspect that the service route steering is not occurring for VPN CAR or CAR SRv6 or the type 2 NLRI for CAR or CAR VPN.

- If (as the shepherd suspects) service route steering is included for type 1 NLRI of CAR, CAR VPN, and CAR SRv6, then section 3 must refer forward to these sections or follow these sections.
- If (as the shepherd suspects) service route steering is also included for type 2 NLRI for CAR, CAR SRv6, and CAR VPN, then must refer forward to these sections or follow these sections.

Issue 1.3: What SRv6 SIDs get passed in CAR SAFI and VPN SAFI?

Section 9.1 differentiates between two distinct use cases for SRv6 (routed service SID and non-routed service SID) for steering SRv6. Section 9.1.1 states regarding the routed service SID:

“An SRv6 locator is assigned for a given intent or color. This locator prefix is distributed using BGP-CAR to ingress PEs in a remote domain. The locator prefix may also be summarized along the path and a summary route distributed to ingress PEs. A new IP Prefix CAR route-type (Type-2) is defined for this purpose.”

Shepherd Questions on the routed use case:

1. Is the routed service SID passed in the type 1 NLRI for CAR SAFI only or the type 1 route for the VPN CAR SAFI or both? (Please list the AFI/SAFI.)
2. When the locator prefix is summarized, does the summary route with the SID get passed in type 2 route for CAR SAFI (1/83 and 2/83), or type 2 route for VPN CAR (1/83 and 2/83) or both?
3. Which one of the examples in section C provide an example of the routed use case?

Section 9.1.2 states regarding the non-routed service SID:

“The SRv6 Service SID allocated by an egress PE is not routed. The service route is advertised by the egress PE with a Color Ext-Comm C.”

The intent-aware path within an egress domain is provided by an SR-TE or similar policy to the egress PE (E, C) [RFC9256]. This (E, C) policy is distributed into the multi-domain network from egress BRs using a BGP-CAR route, towards ingress PEs in other domains.” (see Editorial issue-4)

The (E, C) CAR route is advertised with an SRv6 transport SID allocated from a locator assigned for the intent C. An SR-PCE or local configuration may ensure multiple BRs in the egress domain that originate the (E, C) route advertise the same SRv6 transport SID.” (see Editorial issue-5)

Shepherd Questions:

1. Is the routed service SID passed in the type 1 NLRI for CAR SAFI only or the type 1 route for the VPN CAR SAFI or both?
2. When the locator prefix is summarized, does the summary route with the SID get passed in type 2 route for CAR SAFI (1/83 and 2/83), or type 2 route for VPN CAR (1/83 and 2/83) or both?
3. Do any of the examples in appendix C show this use case?

Shepherd's questions on non-routed use case:

1. Which NRIs is the service route advertised in (AFI/SAFI = 1/1 or 2/1 or 1/128 or 2/128) with the color extended community?
2. Is the non-routed service SID case passing SR-Candidate routes in BGP using SR-TE Policy (AFI/SAFI 1/73, 2/73) [draft-ietf-idr-bgp-segment-routing-te-policy]?
3. What is an SR-PCE?
4. Do any of the examples in section C provide an example of the routed use case?

Action item:

- Clarify the text in sections 9.1, 9.1.1 and 9.1.2 based on shepherd's questions regarding SRv6 SIDs.

Issue 1.4- Why should CAR SAFI and VPN CAR SAFI have type 2 routes?

Question-3.1: If color can be assigned via an SRv6 locator or an LCM, why is it necessary to provide type 2 (IP Prefix) inside the CAR and VPN CAR NLRIs?

The editors of draft-ietf-idr-bgp-car-02 stated in WG LC the reason as:

“But by using a distinct CAR SAFI to distribute these transport routes, we avoid overloading of the IPv6 unicast SAFI (2/1) which also carries Internet service routes, as described in Section 9.3.” (DJ, https://mailarchive.ietf.org/arch/msg/idr/ayRYqImVVro48_AfqYSh4lukBR4/)

Section 10 of draft-ietf-idr-bgp-car-02 states:

“An IP prefix CAR route is a route type that carries a routable IP prefix. It may be originated into BGP CAR SAFI either from an egress PE or from a BR in a domain.

It is used for cases where a unique routable IP prefix is assigned for a given intent or color. Color is not essential to distinguish the CAR route.”

Reminder: LCM-EC conveys intent/color associated with route. When traversing a network domain where a different color is used for next-hop resolution, BGP Color EC may additionally be used as in Section 2.10

Shepherd's questions:

Question 1: What overloading were the CAR authors concerned about?

[private response from authors]: We have clarified in previous exchanges that overloading refers to the mixing of IPv6 Internet routes and SRv6 infrastructure (locator) routes. It is described in the draft too. This is an example of breaking the walled garden, which CAR SAFI provides a way to avoid.

Question 2: Is there a close dependency between the Type 1 Route (E, C) for CAR and the Type 2 Route (E) that requires it to be included as type different types within the same NLRI?

Question 3: If there is a dependency, what procedures define this dependency?

Question 4: Does the inclusion of Type 2 NLRIs provide some ordering benefits?

The authors have pointed to section 9.1.2 for the non-routed case with the following text that provides the link for this dependency:

“ABRs in the egress domain advertise an IP Prefix CAR route for a locator prefix that covers the transport SID allocated by the egress ABR for this (E, C) route. This IP Prefix CAR route (Type-2) is distributed across BGP hops in the underlay towards the ingress PE similar to [the] previous case and may be summarized.”

In private email, the authors have suggested the following definition is contained in section 9.1.2.

“Both Type 1 and Type 2 CAR routes are for intent-aware routing in [a] multi-domain network. Type 1 route is used when [the] same IP address on a router needs to be used to create multiple intent-aware paths. Type 2 route is required when intent and IP Address have a 1:1 relationship. They both use the exact same procedures for path availability, color-based next hop resolution, and LCM when present.” (9/5/2023, DJ)

Shepherd comment: The shepherd feels the text in 9.1.2 does not provide this level of clarity.

In private email, the authors have also suggested the following:

“As described in draft sections 10 and 9.1, there is a 1:1 mapping between the SRv6 locator assignment and intent/color. Hence LCM is not mandatory, since the locator prefix can also be used to map to color on a router based on a local route-policy. But the operational best practice is described to attach [a] LCM so that all procedures can be used automatically.]

In a private email, the authors have suggested:

“There may be many Type-1 routes/SIDs that are dependent on a given locator [summary Type-2] for resolution and reachability?”

Shepherd comment: The shepherd feels the text in 10 does not provide sufficient clarity on what is done by routing policy for CAR type-2 routes.

Question 4: Is this locator prefix a summarization of other transport SIDs? How does this provide an example of dependency to the reader?

Action item:

- Prepare a talk for the 9/25 interim that discuss why
- Improve the text in 9.1, 9.1.2, 9.1.3 and section 10 to include a description of the inter-dependency of type 1 and type 2 routes.

Issue 1.5 – Proof based on section 9.1.1?

DJ in his comment on “no change” in the methodology mentions section 9.1.1 provides a use case of a routed SID based on the following reasoning:

- “As described in Section 9.1.1 for SRv6, the usage for the IP prefix route to resolve the Service SID via an underlay routed locator prefix or summary route (ref. Section 5 of RFC 9252). This is equivalent to the locator prefix being distributed in IGP Flex-Algo.
- Being a regular BGP IPv6 route as specified, it follows the exact semantics as RFC 4271 and RFC 2545, allowing IPv6 routing/forwarding and summarization for these prefixes.
- But by using a distinct CAR SAFI to distribute these transport routes, we avoid overloading of the IPv6 unicast SAFI (2/1) which also carries Internet service routes, as described in Section 9.3. Hope this clarifies.”

RFC9252 states:

“The SRv6 Service SID **SHOULD** be routable (refer to [Section 3.3](#) of [\[RFC8986\]](#)) within the Autonomous System (AS) of the egress PE and serves the dual purpose of providing reachability between ingress PE and egress PE while also encoding the SRv6 Endpoint Behavior.”

Shepherd's comments:

- Based on RFC9252, it is appropriate that the service SID (an IPv6 prefix (2/1)) is routed via a transport location prefix (CAR SAFI (2/83) or VPN CAR (2/84) or a summary route (type 2: IP Prefix). Since the IPv6 AFI (2) with CAR SAFI (2/83) or VPN CAR SAFI (2/84) are NLRIs, it is possible to summarize these prefixes.
- RFC8986 states in section 3.3:
 - “Most often, the node N would advertise IPv6 prefix(es) matching the LOC parts covering its SIDs or shorter-mask prefix. The distribution of these advertisements and [the] calculation of their reachability are specific to the routing protocol and are outside of the scope of this document.
 - This underscores the need for draft-ietf-idr-bgp-car-02 to specify this in the procedures.
- draft-ietf-idr-bgp-car-02 does not have enough specific text (in section 9.1.1 and section 10) that applies that indicates Type-2 routes provide summary routes. The concept of a type 1 and type 2 NLRIs for CAR SAFIs should be introduced in section 1 and procedures pointed to throughout sections 2, 3, 8, 9, and 10.

Action item:

- Rework sections 2, 3, 8, 9, and 10 to provide clarity.

Shepherd Concern: Redistribution of Customer Routers

Reason 1: The “type” field in the EVPN used for the Type-5 type field IRB functions caused problems with interoperability in EVPN. Type 5 allowed a redistribution of customer routes (L3VPN and DC VRF) into EVPN. The method of EVPN Type 5 redistribution caused a breakdown of the walled garden in some implementations under some cases. In some these cases, the incorrect redistribution of these routes caused loops. The key point is that Customer routes need to stay separate from the CAR and VPN CAR SAFIS.

Shepherd Comments: It must be very clear in all sections of the documents that the NLRIs defined in this document (AFI/SAFI 1/83, 2/83, 1/84, and 2/84) are used only for infrastructure routes.

Reason 2: The shepherd notes that the authors and editors of draft-ietf-idr-bgp-car have indicated that the RFC8277 format limits the future flexibility of adding TLVs. However, the tight binding of TLVs to CAR and VPN CAR SAFIs means these are not available for general use outside of the CAR and CAR VPN SAFI.

Action items:

- Discuss on email list and in Talk 9/25 interim
 - why EVPN Type-5 incorrectly redistributed customer routes.
 - why Type-2 does not have this problem (interim and on-mail list).
 - Binding of TLV format tightly to CAR and CAR VPN on list

Issue 1- Action item summary for issues:

Shepherd's action items to address Issue 1.1 – What NLRIs are included

- The abstract should include the fact that two new SAFIs are being created.
- Section 1 paragraph 1 should indicate that this draft specifies two new SAFIs (CAR (83) and CAR VPN (84). This paragraph should clearly state which AFIs these two SAFIs apply to. This should include AFI/SAFIs equal to 1/83, 2/83, 1/84, 2/84.
- Section 2.1 discussion of the database should indicate this document only defines the CAR SAFI for AFI=1 and AFI=2.
- Section 8 should clearly define which AFIs this document defines the VPN CAR SAFI.

Shepherd's actions items for Issue 1.2 – CAR and VPN pass only infrastructure routes.

- Sections 3, 8, 9 and 10 – need to provide a consistent picture of what NLRIs are supported for service routes and infrastructure routes for intent-based/color aware routes.

Shepherd's action items for Issue 1.3 – What SRv6 SIDs get passed in CAR and VPN CAR SAFI

- Clarify the text in sections 9.1, 9.1.1 and 9.1.2 based on shepherd's questions regarding SRv6 SIDs.

Shepherd's action items for Issue 1.4 -

- Discuss on list and prepare a talk for the 9/25 interim that discuss why type-2 routes were added in -02
- Improve the text in 9.1, 9.1.2, 9.1.3 and section 10 to include a description of the inter-dependency of type 1 and type 2 routes.

Shepherd's action items for Issue 1.5 -

- Clarification of text in 9.1, 9.1.1, and 9.1.2 were already noted in issue 1.3 and 1.4

Shepherd's concerns:

- Discuss on email list and in Talk 9/25 interim
 - why EVPN Type-5 incorrectly redistributed customer routes.

Shepherd's Report for CAR WG LC Issues

- why Type-2 does not have this problem (interim and on-mail list).
- Binding of TLV format tightly to CAR and CAR VPN on list

Shepherd's Editorial Comments Related to Issue-1

Editorial Comment 1:

The abstract of draft-ietf-idr-bgp-car-00 states

This document describes a BGP based routing solution to establish end-to-end intent-aware paths across a multi-domain service provider transport network. This solution is called BGP Color-Aware Routing (BGP CAR). (*Editorial Note 1)

Note that “BGP based” should be indicated as “BGP-based.”)

Editorial Comment 2:

Section 2.1 describes the BGP CAR Data model as encoding of the following in the NLRI:

- NLRI Key: IP Prefix, Color
- NLRI non-Key fields: MLS label stack, Label Index, SRv6 SID list, etc. (*2)

Please note that “etc” is an inappropriate inclusion for a specification. Remove it and give full list.

Editorial Comment 3: Section 2.1

Also included in the BGP update packet with issue is:

- BGP Next Hop (*3)
- AIGP Metric (*3)
- Local-Color-Mapping Extended Community (LCM-EC)

[*3 The MP BGP [RFC4760], it is assumed the BGP Next Hop is used, but this is not stated. Editorial changes: BGP Next Hop [RFC4760], AIGP Metric [RFC7311].)

Editorial Comment 4, Section 9.1.2

“The intent-aware path within an egress domain is provided by an SR-TE or similar policy to the egress PE (E, C) [RFC9256]. This (E, C) policy is distributed into the multi-domain network from egress BRs using a BGP-CAR route, towards ingress PEs in other domains.”

SR-TE – needs to be defined. If you mean SR-TE to be draft-ietf-idr-segment-routing-te-policy, then it needs to be referenced. If you mean something specific to RFC9256, please details.

Editorial Comment 4, Section 9.1.2

Text: “The (E, C) CAR route is advertised with an SRv6 transport SID allocated from a locator assigned for the intent C. An SR-PCE or local configuration may ensure multiple BRs in the egress domain that originate the (E, C) route advertise the same SRv6 transport SID.” (see Editorial issue-5)”

The SR-PCE needs to be defined.

CAR-WGLC-Q1-Issue-2: Why Add IP Prefix?

Original Text:

- **[Jeff Haas two Response] to DJ-R-2 – Clarifying IP-Prefix**
 - **Jeff-R1:**
<https://mailarchive.ietf.org/arch/msg/idr/xTpg9qBDRJM7EmAB64E6wjScd-Q/>
 - **Jeff-R2** https://mailarchive.ietf.org/arch/msg/idr/P4hCavFklsH5zQ5y4kMOc4_Xrck/
 - **DJ-Jeff-R1 (DJ-R-4):**
<https://mailarchive.ietf.org/arch/msg/idr/lmqA2VSozVemjx8l0WAol37mXCM/>
 - **DJ-Jeff-R2 (DJ-R-5)** <https://mailarchive.ietf.org/arch/msg/idr/4k8TOj-ag4NmNQm7bAafIkeJNtk/>
 - **Jeff's R1 text:**
 - Per Sue's request, I'll be doing another top to bottom review of the -car draft and then after that re-doing the comparative analysis between the two. I'll be getting to the new additions to -car later on, but this thread provides an opportunity to ask a targeted clarifying question:
 - DJ's text from DJ-R2: "The IP prefix route type is an extension for the specific case where IP Prefix == intent or color, such as in SRv6 where a locator is assigned for a given intent. Here there is no requirement to have multiple instances of the prefix with different colors, hence no need to have the color in the NLRI."
 - This point is understood. Did the authors consider simply designating a targeted color, say 0 on the existing type 1, to avoid creating a new type for the NLRI? Compare vs. the best effort transport-class in the more recent -ct work.
 - **DJ-Jeff-R1-P1:** There are of course multiple options when it comes to encoding.
 - **Jeff-R2-P1:** As we have been discussion for many months.
 - **DJ-Jeff-R1-P2:** As specified in RFC 9256, Color is a non-zero value.
 - **Jeff-R2-P2:** This partly why I offered zero as an example.
 - **DJ-Jeff-R1-P3:** As specified in RFC 9252, steering for a routed SID does not rely on Color-EC. Resolution over best effort paths do not rely on Color-EC.
 - **Jeff-R2-P3:** Understood.
 - **DJ-Jeff-R1-P4:** It did not seem necessary to overload the Color value 0 and redefine all the existing semantics.
 - **Jeff-R3-P4 Similarly understood.**
 - **DJ-Jeff-R1-P5:** The use of a separate route-type is cleaner and allows consistency with existing steering/resolution semantics for service routes. A separate route type for IP prefix also provides a clear distinction for the behavior of this route.
- **Jeff-R3-P 5:** think where I have a level of remaining discomfort is the applicability of the LCM and color extended communities. In the case of the color extended community, we already have a more general sense of how that applies for route types such as BGP-LU. I think it's reasonable to consider this new type as having largely the same existing semantics.
 - **DJ-Jeff-R2-P1:** Yes. This is described in Section 2.5 and elaborated in other sections.

- **Jeff-R3-P6:** LCM, however, was previously introduced to handle the type 1 semantics where the intent color present in the NLRI required a local override based on crossing color domains. It signals that intent resolution is to proceed over the LCM value rather than the NLRI tag color.
 - **DJ-Jeff-R2-P2:** Yes. Usage described in Section 2.9.3 and 2.10.
- **Jeff-R3-P7:** The documentation in section 10 begins by pointing out that "color is not essential", and then provides some use cases. It then moves to "color is still useful" situations.
 - **DJ-Jeff-R2-P3:** To be precise, it says color is not essential to distinguish the route (as I've also clarified in this thread); but it is useful as an indicator of the NLRI color for other purposes as listed in the section.
- **Jeff-R3-P8:** When LCM is present, this new type is - I believe - functionally equivalent to the type 1 -car NLRI. Correct?
 - **DJ-Jeff-R2:** Yes. The LCM-EC usage (as well as Color-EC) is the same for both route types, as described/referenced in Section 10.
- **Jeff-P9:** If so, it might be worth some text recommending that when color is semantically important to use type 1. If not, some additional text covering how they are functionally or semantically different than type 1 entries would be helpful.
 - **DJ-Jeff-R2:** We used the approach of describing when the IP prefix route is applicable. But we can try to add clarifying text.

Shepherd's action items:

- Better description of applicability of LCM in section 2.5, 2.9.3, 2.10, 8, 9, and 10.
- Consistent story on color in section 10 which includes (type 2 + color) = Type 1.
- Add text on when color is semantically important to Type 1

CAR-WGLC-Q1-Issue-3 – Clarity of Support for IP Prefix in AFI/SAFI Section 10 + No Implementations

Original Text:

- **Kalraj-R-1 response to DJ – Clarity on AFI/SAFI Support + Section 10**
 - few meta comments, and a couple questions follow.
 - **link:** <https://mailarchive.ietf.org/arch/msg/idr/UUmdgjl6XzZsS7SZZMuFzvShgXI/>
 - **[Kaliraj-R1-Part1]** regarding DJ's comment:
 - This route type and it's uses are also applicable to the VPN layer, hence to VPN CAR. Traditional IP prefix advertisement, such as BGP IPv6 or BGP-LU (<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-car-02#section-10>)
 - So, CAR is subsuming other BGP families (VPN, IPv6-Unicast, LU) into itself? But without specifying clearly how the business logic of those families will work inside CAR. All procedures specified in car draft are for E,C NLRI.
 - **[DJ-R-3 response to Kaliraj-R-1]:** here is no change in scope compared to previous versions of the CAR draft. Since there is a new route-type, it's applicability is specified where necessary. For instance, it does apply to the PE-CE peering similar to the (E,C) route, hence it's specified as above.
 - Link for DJ-R-3: <https://mailarchive.ietf.org/arch/msg/idr/ZEQyDYc0nS19PR8KAs2v1J12PkI/>
 - **[Kaliraj-R1-Part2]:** There is no clarity on whether CAR is a Service family or Transport family, or the procedures thereof. e.g., VPN-CAR (<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-car-02#section-10.2>) has no procedures specified.
 - **[DJ-R-3 response to Kaliraj-R1-Part 2]:** The VPN CAR semantics and procedures are same as already specified in section 8, this section is just specifying the new route-type. But we can explicitly state it too. The use-case as before is to establish an intent-aware path across multiple domains over which service routes (and traffic) can be steered over. These paths can extend not just PE-PE but also extend to customer networks. This is described in the problem statement draft:
 - **[Kaliraj-R1-Part3]:** Also, like Moshiko points out below, there is confusion on whether Color should be in the NLRI or not.
 - See <https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-car-02#section-2.9.2>
 - And <https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-car-02#section-10.1>
 - Sec 10.1, removed Color from Type 2 NLRI, and is left with no Distinguisher to avoid path-hiding in path-selection pinch points.
 - **[DJ-R3-response to Kaliraj-R1-Part3]:** I've responded to Moshiko's message. But to reiterate,
 - The IP prefix route type is an extension for the specific case where IP Prefix == intent or color, such as in SRv6 where a locator DR# is assigned for a given intent.
 - Here there is no requirement to have multiple instances of the prefix with different colors, hence no need to have the color in the NLRI
 - **[Kalirja-R1-Part4]:** Meta comments:

- Things seem to be done in a hurry, and got added ad-hoc, with the original problem of intent-driven service mapping lost.
- I think there is not enough clarity in the document that we can do a more detailed review wrt Part 3 questions. But we're trying.
- **[DJ-R3-response to Kaliraj-R1-Part4]:** Discussion on comment on the number of CT revisions versus CAR.
 - **[Editor note:** The details this comment have been since the WG LC specifically requested no comparisons.]
- **[Kaliraj-R1-Part5]:** Also, as per the Implementation report:
 - <https://wiki.ietf.org/group/idr/implementations/draft-ietf-idr-bgp-car>
 - Following constructs have no implementation: VPN CAR, Type 2 Prefix Route-type, LCM EC
 - **[DJ-R3 response to Kaliraj-R1-Part5]:**
 - [Editor's comment – I did not find an item that directly responded to this comment. I may have missed it.]
- **[Kaliraj-R1-Part6]:** Meta comment: These things bring me to ask the following questions for the WG, to save everyone's time:
 - Are we rushing towards a LC for a completely new set of extensions in car-02 with no/incomplete procedures?
 - Should we go back to the drawing board, and reset the LC (for CAR only)?
 - Do we really want a family that subsumes other existing families (VPN, IPv4-Unicast, IPv6-Unicast, LU)? [Are we] considering the security implications and filtering problems that come with it? There is no walled garden anymore.
 - **[Editor's comment:** The walled garden issue is further discussed in Issue-4 on security.]
 - **[DJ-R3-response:** I've tried to respond to technical comments. It's worth restating that the semantics and procedures used by the CAR SAFI are consistently aligned with existing SAFIs such as BGP-LU and BGP IPv6, and the draft has described the specific extensions as needed.
 - **[Editor's comment:** DJ's comments on changes in the CT document were removed. The forum rules prohibited comparisons.]
- **Robert-R-1: [Reply to Kaliraj]:** Please elaborate how CAR "subsumes other existing families (VPN, IPv4-Unicast, IPv6-Unicast, LU)"? Having color marking applicable to next hop(s) does not make any absorption of service address families which happen to use such next hop(s).
 - **Link:**
https://mailarchive.ietf.org/arch/msg/idr/UuNDOE8nGK_fIRP5g_LbUgCs7Ok/
- **Natrajan (nats-R-1): [Reply to Robert-R-1]** We have specific AIs from the WG chairs to elaborate on the new extensions and new NLRI types that have been introduced in CAR on 7/7/2023. There are specific email threads addressed by the workgroup chairs towards me (Nats) to reply towards the same. We will inform you once we are done with that.

- In the meantime [your comment: “Having color marking applicable to next hop(s) does not make any absorption of service address families which happen to use such next hop(s).”
 - I am unable to understand the context in which you are making the above statement. Could you please clarify?
 - Could you point to sections in the CAR draft (like Kaliraj did in his email), so that we can understand your statement? If you are not pointing to anything in the CAR draft, it would be difficult to reply
- Link: https://mailarchive.ietf.org/arch/msg/idr/g_2MDTiB_bPQb0gg6rH8-f8PA3w/

Shepherd's comments

All of these issues reinforce issues 1.1 to 1.5

Shepherd's action items:

- Add a clear description of AFI/SAFI included in CAR SAFI, VPN CAR SAFI (mentioned in issue 1-1, issues 1-2)
- Modify the security section 13 to include both CAR and CAR VPN to include the following issues:
 - Any better/worse security with tightly bound type 1 and type 2 routes. (perhaps the last paragraph).
 - Pointing to clear policy about service routes versus infrastructure routes.

CAR-WGLC-Q1-Issue-5 – Lack of Clarity for Procedures in Sections 8-10

Original Text:

Sue-R-2: (part-1) – Lack of Clarity on AFI/SAFI support

Link: https://mailarchive.ietf.org/arch/msg/idr/YpbOiWdz0qE_KvRCbJG-Z5X2i4U/

Would it be valid to restate your reason for the CAR AFI/SAFI as: “The CAR SAFI is intended to update the BGP transport functions in BGP-LU (applicable to just MPLS) for new transports (such as SRv6). It maintains the exact flat routing semantics of BGP-LU/IP without any need for VPN-like import.”

DJ-R-6: Sure – that would be fine.

Link: https://mailarchive.ietf.org/arch/msg/idr/jc2kkkKNGoP439aWISOa94kX_1w/

- **Kaliraj-R-1 comment on DJ-Response-2:**
https://mailarchive.ietf.org/arch/msg/idr/8i2mWlg3YTxY4RN1zjcD0_pXyGQ/
 - **[Kaliraj-R-1]:** Here you seem to agree with Sue that families using CAR SAFI are at transport-layer. But then, you go on to say:
 - DJ-quoted: “BGP CAR can be enabled between PE-CE. To advertise the customer’s CAR routes between PEs, VPN CAR is required.”
 - **[Kaliraj-R-1]:** Do you see this usage of CAR and VPN-CAR families as a transport-layer of service-layer usage?
 - **[Kaliraj-R-1]:** Also, section 8 or 10 lack procedures on how this usage interacts with existing AFI/SAFIs (e.g., 1/1 or 1/128) that are deployed today on PE-CE or PE-PE sessions?
 - **[Kaliraj-R-1]:** How would you state your reason for the CAR VPN AFI/SAFI? (This AFI/SAFI contains a spot for the RDs) see section 10.2. Why support RD?
- **Sue-R-2: (part-2):** How would you state your reason for the CAR VPN AFI/SAFI? (This AFI/SAFI contains a spot for the RDs) see section 10.2. Why support RD?
 - **DJ-R-6:** As stated below and for the use-cases in draft-hr-spring-intentaware-routing-using-color-02, BGP CAR can be enabled between PE-CE. To advertise the customer’s CAR routes between PEs, VPN CAR is required. The semantics are as per RFC 4364, the RD uniquely distinguishes CAR routes from different customers (VRFs).
 - **DJ-R-3:** This is already described in section 8 of the CAR draft. It’s not new, was present in earlier revisions and discussed during adoption.
- **Sue-R-2 (part-3):** You have two NLRI types: Color-Aware Route NLRI, IP-Prefix NLRI. The color-aware route NLRI Intent (E,C) has been discussed for 2 years in IDR and SPRING. The concept of Intent has been defined in the IRTF. The IP-Prefix NLRI is new.
 - **DJ-R-3:** It may be new in the draft but semantically its same as an BGP IPv4/BGP IPv6/LU IP prefix. Hence not a new behavior.
 - **DJ-R-3:** The key difference is operational – as the CAR SAFI enables a clear separation between transport routes and service routes by not overloading the SAFIs that also carry service routes such as SAFI 1 or 4 - as already stated in Section 9.3.
- **Sue-R-2 (part-4)** You state: “The IP prefix route type is an extension for the specific case where IP Prefix == intent or color, such as in SRv6 where a locator is assigned for a given intent. Here there is no requirement to have multiple instances of the prefix with different colors.

- (For WG Note: Jeff Haas already asked about whether the color = 0 would have been useful.
- https://mailarchive.ietf.org/arch/msg/idr/P4hCavFklsH5zQ5y4kMOc4_Xrck/
- <https://mailarchive.ietf.org/arch/msg/idr/xTpg9qBDRJM7EmAB64E6wjScd-Q/>
- **DJ-R-2:** I have responded to Jeff. Please let me know if you have any additional concerns.

Shepherd's comments:

1. If BGP CAR is enabled between PE-CE, are the routes being pass infrastructure routes or service routes?

DJ states:

“DJ-R-6: As stated below and for the use-cases in draft-hr-spring-intentaware-routing-using-color-02, BGP CAR can be enabled between PE-CE. To advertise the customer's CAR routes between PEs, VPN CAR is required. The semantics are as per RFC 4364, the RD uniquely distinguishes CAR routes from different customers (VRFs).”

2. If BGP CAR is passing customer routes you are entering into the space of Type-5 EVPN. Please clarify. RD support infrastructure routes that link customer spaces as well as

Shepherd's action items:

- Inform the shepherd immediately if you are passing customer routes in CAR or VPN CAR SAFIs in type 1 or type 2 routes. Send an example of the PE-CE interchange.
- Clear define in sections 2, 3, 8, 9, and 10 – what is a infrastructure route, a service route, and a customer route.

CAR-WGLC-Q1-Issue-6: Relationship among multiple drafts related to SRV6

Original text:

- **Sue-R-2 (part-5):** draft-wang-idr-cpr-02 gives some deployment insights from a set of operators and vendors. Did insights from this draft cause you to add the IP-Prefix? If so, which insights or customer deployments?
 - **DJ-R-2:** Not at all. The precedence for distributing SRv6 locator prefixes in routing is well established in IGP (IGP Flex-Algo - RFC 9350) and BGP (<https://datatracker.ietf.org/doc/draft-agrawal-spring-srv6-mpls-interworking/>)
 - **DJ-R-2:** It is a natural outcome of the definition of the routed SID (<https://datatracker.ietf.org/doc/html/rfc8986#name-sid-reachability>) and its usage for steering is described in RFC 9252 (<https://www.rfc-editor.org/rfc/rfc9252.html#name-bgp-based-l3-service-over-s>)
 - **DJ-R-2:** BGP CAR just follows the established model. This was one of the use-cases for defining a route-type.
- **Sue-R-2 (part-6):** Thanks for helping the WG understand the intent of your design.
 - DJ-R-2: DR# Hope the above clarifications are helpful.

Action item:

- Clarify the relationship between draft-ietf-idr-bgp-car-02 and draft-wang-idr-cpr-02 during the IDR interim on 9/25.

CAR-WGLC-Q1-Issue-7 – Using CAR SAFI and CAR VPN SAFI for transport

Original Text:

- Robert Raszuk: Comment 2 on CAR/SAFI
- Link: https://mailarchive.ietf.org/arch/msg/idr/gKBY651k2da_3qtNNhWExrqfnoA/
- Text: address to Kaliraj & Sue:
 - **Robert-R2-P1:** As I mentioned in my note my understanding is that CAR routes augment transport so [these routes] are exclusive to next hop's color marking. You still need SAFI 1/1 and 1/128 or 2/1 and CAR SAFI would be just carrying colored next hops for those service SAFIs. And CAR will not in any form or shape replace those. Interaction between services SAFIs and CAR SAFI is via next hop resolution.
 - [Editor's comment (to Robert)]
 - CAR SAFI – per section 2.9 is 1/83 or 2/83 paragraph
 - Has 2 forms:
 - CAR NLRI (color, prefix, options)
 - IP Prefix (prefix, color, options)
 - Use: augments transport
 - CAR VPN SAFI: Per section 8 does not indicate which AFIs it is valid for.
 - Has two forms
 - CAR IP prefix (color, prefix, options)
 - IP Prefix (prefix, color, options)
 - **Robert-R2-P12:** That is why I asked in the former note where your assumption came from. Your quotes to some lines of the CAR draft are simply out of context as you need to put them into overall CAR intention and keep in mind that CAR is only transport augmentation.
 - [Editors comment]
 - If CAR is only: transport function, then why does the encoding matter:
 - It appears the IP Prefix has taken over the CT function.

Shepherd's action items:

- draft-ietf-idr-bgp-car – should clarify the interaction between NLRI's with 1/1, 1/128, and 2/1 and the CAR SAFI and VPN CAR SAFI.

CAR-WGLC-Issue-8- Use of CAR-SAFIs for Services Families

Original text:

Kaliraj-Response-2:

- Link: https://mailarchive.ietf.org/arch/msg/idr/vxkY99dcOcNoW6-3G0dRzbJH_1U/
- Comment: Text in car-02 draft indicating CAR families usage as Service families:
 - <https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-car-02#section-8>
 - BGP CAR is enabled between CE1-PE1 and PE2-CE2
 - BGP VPN CAR is enabled between PE1 and PE2
 - In above text, CAR is used as Service family between CE-PE, where usually AFI/SAFI 1/1 is used. and VPN-CAR is used as Service family between PEs, where usually AFI/SAFI 1/128 is used
 - <https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-car-02#section-10>
 - “A couple of applicable use-cases:
 - SRv6 locator prefix,
 - Traditional IP prefix advertisement, such as BGP IPv6 or BGP-LU
 - Presuming "BGP IPv6" means AFI/SAFI 2/1, this text says CAR Type2 route is used to carry IPv6-Unicast service routes.
- **DJ- Response to Kaliraj-2:**
 - Link: <https://mailarchive.ietf.org/arch/msg/idr/W6Z24OkxXUOcUUZdlq5YG2m7MUK/>
 - Text : A clarification to your comments below - CAR SAFI is not being used instead of SAFI 1/1.
 - SAFIs 1/1 and 2/1 continue to be used for existing customer service routes and the resulting VPN routes (eg. in SAFI 128) will be steered along the provider network CAR paths as specified in section 3.
 - CAR SAFI can be independently and optionally enabled between CE-PE when a customer is running CAR within their network and have a need to extend it across the provider network. For example, for a CsC or equivalent scenario that needs intent-awareness.
 - There isn't a specific procedure to be described, as the SAFIs are orthogonal to each other. CAR SAFI routes on the PE will be advertised via VPN CAR SAFI as described in Section 8. We can add some more clarifying text if that helps.

Shepherd's Action items:

(note these action items repeat items from issue 6 and 7)

- Inform the shepherd immediately if you are passing customer routes in CAR or VPN CAR SAFIs in type 1 or type 2 routes. Send an example of the PE-CE interchange.
- Clear define in sections 2, 3, 8, 9, and 10 – what is a infrastructure route, a service route, and a customer route.

OPS-DIR review

Link: <https://mailarchive.ietf.org/arch/msg/idr/4SgOciKKtA9nap1M-f6eHRKNtWc/>

Reviewer: Yingzhen Qu
Review result: Has Issues

Hi,

Thanks for the draft.

I am the assigned OPSDIR reviewer to conduct an "early" review of this draft.

General comments:

There are lots of abbreviations in the draft. I'd suggest to add them in the terminology section. For example, I'd assume BR means Border Router, but there might be different guessing.

In this draft, it says E is globally unique, which makes E-C in that order unique. Can you please explain a bit more about the second unique? I suppose it's possible to have two different source nodes, E1 and E2, all reach destination E with color C, correct?

The draft has an informative reference to [I-D.hr-spring-intentaware-routing-using-color], which is an important problem statement for this solution. Will the problem statement draft progress as well? Even so, to improve the readability of the bgp-car draft, I'd suggest adding some text for a brief introduction of the problem.

IP Prefix NLRI was added in version -02. The use case is where a unique routable IP prefix is assigned a given intent or color. In other words, the IP is overloaded with a color. The same can be achieved using an IP with a color. I'm not totally convinced that this type 2 NLRI is needed. Please clearly specify when it should be used.

Please consider adding a section for operation considerations. There are pieces of information about operation and deployment scattered in the document, please consider group them together.

There are quite some sentences missing "." at the end. Please do an editorial pass and fix them.

Detailed comments with line # from idnits:

478 The value set (or appropriately incremented) in the AIGP TLV
479 corresponds to the metric associated with the underlying intent of
480 the color. For example, when the color is associated with a low-
481 latency path, the metric value is set based on the delay metric.

483 Information regarding the metric type used by the underlying intra-
484 domain mechanism can also be set.

Date 9/11/2023

Shepherd's Report for CAR WG LC Issues

comment: This statement lacks a clear definition how the metric should be set.

486 If BGP CAR routes traverse across a discontinuity in the transport
 487 path for a given intent, add a penalty in accumulated IGP metric
 488 (value by user policy). For instance, when color C1 path is not
 489 available, and route resolves via color C2 path (e.g., Appendix A.3).

How about the case where encapsulations are different? For example, SR policy in one AS and IGP-FlexAlgo in the other AS vs. SR Policy in both ASes.

Section 2.7

504 The (E, C) route inherently provides availability of redundant paths
 505 at every hop, identical to BGP-LU or BGP IP.

"every hop" is a bit confusing here since it may mean an IGP hop within an AS. To my understanding, this section means ECMP or backup paths can provide protection in case of failure within an AS domain without impact other ASes.

"Path Availability" as the section title is not very clear. How about something like "Inherent Path Protection"?

513 BGP ADD-PATH should be enabled for BGP CAR to signal multiple next
 514 hops through a transport RR.

I'd suggest to change to "SHOULD be enabled".

526 The BGP CAR solution seamlessly supports this (rare) scenario while

I'd suggest adding a small paragraph explaining why this is a rare but useful case. I would guess the two domains used to belong to different administrators, now they're trying to merge under one admin domain. nits: personally I don't like how "(rare)" with parentheses is used here, but I'd leave this to the authors.

806 NLRI instead of the BGP Prefix SID attribute. The BGP Prefix SID
 807 Attribute SHOULD be omitted from the labeled color-aware routes when
 808 the attribute is being used to only convey the Label Index TLV.
 Add a reference to Appendix D?

848 BGP CAR SRv6 SID TLV definitions provide the following benefits:

850 * Native encoding of SIDs avoids robustness issue caused by
 851 overloading of MPLS label fields.

853 * Simple encoding to signal Unique SIDs (non-transposition),
 854 maintaining BGP update prefix packing

856 * Highly efficient transposition scheme (12-14 bytes saved per
 857 NLRI), also maintaining BGP update prefix packing
 minor: I don't think the text belongs to the encoding section. Maybe part of "Operation Considerations"?

1019 * If multiple instances of same type are encountered, all but the

Shepherd's Report for CAR WG LC Issues

1020 first instance MUST be ignored.

1022 * If multiple instances of same type are encountered, all but the
1023 first instance MUST be ignored.

nits: please remove the repetition.

1025 * A TLV is not considered malformed because of failing any semantic
1026 validation of its Value field.

Q: When should a TLV be considered malformed? and how should it be handled?

1033 3. Service route Automated Steering on Color-Aware path

nits: Service Route Automated Steering on Color-Aware Path

Please check to make sure all section titles are consistent.

1044 destination, per-flow, CO-only. For brevity, in this revision, we

1045 refer the reader to the [RFC9256] text.

nits: maybe change to something like "For brevity, please refer to [RFC9256]
section X for detail."?

1047 Salient property: Seamless integration of BGP CAR and SR Policy.

minor: personally I don't think this sentence belong to this section.

1055 4. Intents

The section title and content don't seem to match. I don't quite understand
the purpose of this section.

1085 A separate document will analyze the BGP CAR supports for 3, 5 and 6.
Any reference?

1097 5.1. (E, C) Subscription and Filtering

Q: how is this subscription sent between routers?

1115 * If A does not have (E2, C1), it will advertise F (E2, C1) to its

1116 peer B

I suppose it meant to be "If A does not have subscription of (E2, C1)"

1124 On-demand filtering procedures are outside the scope of this

1125 document.

what's "on-demand filtering"?

1138 Two key principles used to address the scaling requirements are a

1139 hierarchical network and routing design, and on-demand route

1140 subscription and filtering.

Q: on-demand filtering is claimed to be out of the scope. (line #1124)

1342 Note: E1 does not need the BGP CAR (451, C1) route

Q: what's the benefit?

1545 7. Routing Convergence

comments: Maybe section 2.7 and 7 should be put together somehow, but I'll
leave this to the authors.

1602 +-----+

Shepherd's Report for CAR WG LC Issues

1603 | NLRI Length | Key Length | NLRI Type |Prefix Length |
 1604 +-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
 Q: what's value of the NLRI Type here?

1664 service route is advertised by the egress PE with a Color Ext-Comm C.
 nits: "Color Ext-Comm" is only used here once while "Color Extended-Community"
 elsewhere.

Nits: this section 9.2.1 and 9.2.2 needs some editorial work. For
 example, each bullet point should be a sentence finished with a ".".

1796 CAR SAFI may also be used for best-effort routes in addition to
 1797 intent-aware routes.

Q: Should a color be specified here? or use the IP Prefix NLRI Type 2?

2008 This extension defines a new SAFI within a BGP and therefore does not
 it should be two new SAFIs: BGP CAR/83 and BGP VPN CAR/84

2307 The examples use MPLS/SR for the transport data plane. Scenarios
 2308 specific to other encapsulations will be added in subsequent
 2309 versions.
 nits: this should be removed.

Thanks,
 Yingzhen

Shepherd's action items:

- 1) Make sure all abbreviations have explanation in introduction,
- 2) Explain In this draft, it says E is globally unique, which makes E-C in that order unique. Explain this concept in section 1 and point back to in section 2.1.
- 3) Explain why two types of NLRI are needed in section 1 or 2.1
- 4) Enhance the explanations on Type 2 NLRI (IP Prefix) to indicate the routes only cover infrastructure routes.
- 5) Add an "operations" or "operational" section
- 6) Fix all editorial issues.