

CAR-CT Adoption Poll – Part 2: Adoption Detailed Summary

This begins a 2-week WG Adoption call (7/6/2022 to 7/20/2022) for the following drafts:

- draft-dskc-bess-bgp-car-05.txt
(<https://datatracker.ietf.org/doc/draft-dskc-bess-bgp-car/>)
- draft-kaliraj-idr-bgp-classful-transport-planes-17.txt
(<https://datatracker.ietf.org/doc/draft-kaliraj-idr-bgp-classful-transport-planes/>)
- draft-kaliraj-bess-bgp-sig-private-mpls-labels-04
(<https://datatracker.ietf.org/doc/draft-kaliraj-bess-bgp-sig-private-mpls-labels/>)

CAR and CT [outside of co-authors]:

[Robert Razsuk] Watching this debate I think the best option is to accept both proposals as experimental. Let the market decide not the mailing list then move one of them to Standards Track and the other one to Historic.

[Jim Uttaro]: I support the adoption of both BGP CAR and BGP CT. I agree [With Robert]. Currently there are two solutions although they are “functionally identical” there are important reasons why operators may opt for one or the other. I am a bit confused as to the driver “We need one interoperable solution”. Kompella and EVPN both can be used to create PWs, Multicast can be deployed using ingress replication, Rosen etc.... The time to drive towards one interoperable solution is before the technologies are specified. I look forward to a discussion of the use cases addressed in order to better inform architect/designers as to which approach is most appropriate for them. Letting the market decide will provide needed perspective to resolve this in favor of one or both specifications.

CAR [outside of co-authors]:

[Zafar] I support adoption of the BGP CAR Draft.

A BGP transport routing solution should maintain the simplicity of BGP IP and BGP-LU. Use of VPN RDs, import/export are overkill for transport nodes and tax for the BGP control plane. BGP-CAR is identical to the BGP-LU model, extended by color.

- BGP-CAR seamlessly works with the SR-TE architecture using Color Extended-Community in service routes and Color in the transport, thus simplifying networks that deploy both BGP and SR-TE and require co-existence and interworking. BGP-CT unnecessarily introduces different mechanisms like mapping communities, creating additional operational overhead and interop issues.
- Supporting signaling of different encapsulations efficiently with extendibility is very important, as we have not just MPLS, but SR-MPLS, SRv6 also being deployed in the same network. The fact that BGP-CT uses the same encoding as BGP-LU is a big limitation.
- Following the discussions during the IDR sessions and the list, the CT solution for multipathing appears to have basic control and forwarding design issues – if it needs to signal multiple routes for the same LSP. What is the purpose of flooding this duplicate information to all nodes ?

NOTE: BGP CAR already has implementations by multiple vendors, with successful interop.

CAR Co-authors:

[Dhananjaya, CAR co-author]: As co-author, I support the adoption of BGP-CAR and oppose CT. His reasons specified in the following email.

https://mailarchive.ietf.org/arch/msg/idr/R-ockSQHvDojiaEu_OgLbDoxBk/

[Swadesh, CAR co-author]: I support adoption of BGP CAR (draft-dskc-bess-bgp-car-05.txt) as a co-author and oppose adoption of BGP CT for the protocol reasons already listed by Dhananjaya Rao (see above) plus four reasons listed in:

<https://mailarchive.ietf.org/arch/msg/idr/PKkqHGQYlrU2or3Af3ABbODq24I/>

[Bruno, CAR co-author]: I support the adoption the CAR solution. IMO, color is part of the NLRI: one need one path per (color, IP destination) just like without color, one want one path per IP destination. So (color,IP) is the Network Layer Reachability Information,

[Luay, CAR co-authors] I support the adoption of the BGP CAR solution. I don't want to repeat what others have already mentioned (co-authors & supporters). I just want to add that as the technologist & architect for most of Verizon's IP networks (Backbone, Internet, L2/L3 VPN services, Wireless Transport/backhaul, IP network slicing) BGP-CAR fits better with our use cases and operational model This should help with making assumptions on what works best for operators, since operators can be different.

Oppose CAR: (2 CT co-authors)

[Kaliraj Vairavakkalai, CT-co-author]: see 6 reasons in

<https://mailarchive.ietf.org/arch/msg/idr/2jkdJD4AP1pyOI7FwslA3KM8aZo/>

[Natrajan Venkataraman, CT-co-author] agrees with Kaliraj and adds reasons for BGP-CAR for customers and developers in:

https://mailarchive.ietf.org/arch/msg/idr/0ZOHIDVii0MhlouedA0Q7EY0_eo/

Approve CT [outside co-authors]

[Tomas Szewczk]: For about 10 years the seamless MPLS services are implemented in Polish NREN (PIONIER). Additionally, the similar service architecture is used among European NRENs connected to GEANT network. Recently we had discussion about extending functionality of our interdomain MPLS services. The BGP-CT proposal seems to answer our functional needs and it seems due to its properties/architecture (similar to most BGP VPN technologies) can be easily and quickly adopted by multiple NRENs.

[Moshiko]: I support adoption of BGP-CT solution. As mentioned in Part 1, I could easily implement BGP-CT in brownfield lab and demonstrated it to my customers. I received positive feedback from different customers for different type of use case metro use cases and across continent infrastructure. It solves both Intra and Inter AS use cases.

[Aravind]: I did try out BGP-CT on the newer Junos versions and it works as expected. I received positive feedback from my customers on the POC conducted for various use cases they had (traffic

redirection, data Sovereignty etc). The deployment aspect being easy to convert an LU based network to being transport aware and having the debuggability aspects much easier is compelling. I support the BGP-CT draft as it solves both green field and brownfield network use cases and keeping the transport agnostic.

[Reshma Das]]: I am writing this email to provide a viewpoint from someone who read and implemented and is currently helping in customer engagement of BGP-CT for JUNOS.

- Firstly reading the draft and understanding the technology was very simple. The concepts and the machinery used is similar to BGP VPN technologies that [are] widely deployed and understood, which makes the time to familiarize BGP-CT short.
- From an implementation standpoint, we were able to extend and re-use time-tested design of L3VPN procedures. We found ourselves in lesser brainstorming sessions about protocol intricacies and could focus our energy in further automation of the implementation that eventually customers liked.

Thereby, implementing BGP-CT was less tedious and error prone as it allows re-use of concepts and procedures that have stood the test of time. Currently talking and working with customers, we get a positive feedback as it solves all their use cases and also it works with various modules and features within BGP that is currently deployed in field (RTC, protection etc).

[Shawn Zhang]: As an end customer, I support the adoption of BGP-CT. Why do I support CT? When deploying inter-as tunnels before, we have hit the problems as:

- .A domain may have intra-AS tunnels with varying TE characteristics (gold, silver, bronze) denoting different SLAs,
- There could be multiple tunnels to the same destination. And different tunneling protocols creating those tunnels in different adjacent domains.. The tunnels may need to be extended inter-domain, while preserving their TE characteristics (aka 'color') end-to-end. Extend BGP protocol to signal these pieces of information.
- Provide inter-op between different tunneling protocols that may be in use in different domains.
- Provide ability to map a Service route to:
 - Tunnels of a certain TE characteristic, with fallback to best-effort tunnels (Default).
 - Tunnels of a certain TE characteristic, without any fallback.
 - Tunnels of a certain TE characteristic, with fallback to [a] different TE characteristics,
- Restrict best-effort service-routes from using 'colored' tunnels. They use only best-effort tunnels.

Mapping to Inter-domain tunnels should work the same way as mapping to Intra-domain tunnels. We have benchmarkingly tested the arch. base on Junos, basically meets our requirements.

[Shawn Zhang] (continued): How BGP-CT can solve the problems? [I appreciate that it does]

- It Implement a new construct viz. "Transport Class", that collects tunnels of a certain TE characteristic.
- The "Transport Class" maps to the "Transport Topology Slice" in 5G Network-slicing .
- A Transport class is identified by a 32bit Color. This Color value aligns with the Color carried in SR transport-protocols like SRTE, ISIS-FlexAlgo.
- And it is carried to neighboring AS-domains as "Transport Route Target" using a new BGP family: BGP-CT
- Implement support in Transport protocols to associate a tunnel to a specific Transport class. RSVP, ISIS-FlexAlgo are supported.

- Service routes carry ‘mapping-community’ (e.g. Color extended community) that indicate the desired SLA. This lets them resolve over tunnels in associated transport-class RIB.
- Implement “Resolution Scheme” construct to provide the more sophisticate fallback schemes. By default, service-routes binding to a transport class use best-effort tunnels as fallback. As you know, fallback is always what we have to consider all time.
- BGP protocol extensions extend the Classful Transport architecture to multi-domain.
- It defines a new Transport Layer BGP family, viz. “BGP CT”, which use [AFI1/SAFIs: 1/76, 2/76]. A new route-target “Transport-Class Route Target” is defined to carry the Transport class ID (Color).
- This new family follows time tested VPN mechanisms (RFC-4364) to leak transport-routes to appropriate Transport Class RIBs. It follows RFC-8277 NLRI encoding. by default, i.e. does not advertise anything to EBGp peers without explicit export policy. The new transport-family uses sane defaults like per-prefix-label,
- [Daisuke]: I support CT. I could create end-to-end path which consists of various underlay path (Flex-algo, RSVP,SR-TE) in Junos. I think CT approach is easy to understand for operators and operators can have flexibility for path selection by resolution scheme concept.

[Daisuke Sugahara]: I support CT. I could create end-to-end path which consists of various underlay path (Flex-algo, RSVP,SR-TE) in Junos. I think CT approach is easy to understand for operators and operators can have flexibility for path selection by resolution scheme concept.

[Gokhan Gumus]: I had a chance to look at the drafts and support adoption of the BGP-CT solution. The concepts are easy to understand and operate. Implementing BGP-CT is easy and can be tested in the lab in a matter of minutes. This solution may help euNetworks to provide end-to-end coloring with the current brownfield deployment instead of moving to SR. We could offer this to our customers without a need to redesign the network. Deployment in brownfield without downtime and no interruption to service traffic. Easy to troubleshoot using RD and RT show commands to follow the color mapping I also like that it is service agnostic which means I can use it for RSVP-TE or SR-TE.

[Gokhan Gumus (continued)]: (in response to Robert on SR deployment) We are not moving toward SR as of today, this may change in the future if SR will bring us something we feel is needed. We may also face SR if we acquire a third party with SR core. Implementing BGP-CT means we can interact with SR without changing behaviour as inter-op is important for my company.

Approve CT: Co-authors –

[Note: The authors comparison given as their technical opinions. For a detailed discussion of the technical opinions, please see the mail thread: Part 3 of the CAR/CT Adoption call at

<https://mailarchive.ietf.org/arch/msg/idr/3RH-Ra4igU4y0sVNC0WM4W0WM0Q/>.

This thread has specific posting requirements.]

[Kaliraj, CT co-author]: As one of the authors, I support adoption of the BGP-CT draft, and I don’t support adoption of the CAR draft. See <https://mailarchive.ietf.org/arch/msg/idr/2jkdJD4AP1pyOI7FwslA3KM8aZo/>

[Natrajan (Nats), CT co-author]: As a co-author, I support adoption of the BGP-CT draft. I DO NOT support the adoption of BGP-CAR for the same reasons that Kaliraj has pointed out below, and I am in

total agreement with that. Nats goes on to specify reasons as a customer and a developer in the following message:

https://mailarchive.ietf.org/arch/msg/idr/OZOHIDVii0MhlouedA0Q7EY0_eo/

[Israel Means, CT co-author]: I endorse the adoption of CT.

- While both CAR and CT address color as a reflection of intent the operational approaches are distinct. With CT existing network constructs such as RD, RT and RTC are leveraged reducing the operational learning curve for MPLS VPN network operators.
- We also find that CT simplifies operations of network structures that include SR and non-SR implementations under a single administrative authority. CT provides an operationally simple approach that facilitates the creation of seamless end-to-end intent-based paths over networks with diverse protocols (example: SR vs RVSP-TE).

[Gyan, co-author]: As co-author of CT, I support IDR WG adoption of CT. His reasons are specified in the following email: <https://mailarchive.ietf.org/arch/msg/idr/S39ZABGn4Ny0khKhWkIbgcENmhs/>

[Deepak Gowda, CT co-author]: BGP CT concepts are much [very] similar to BGP L3VPN and the existing BGP VPN infra structure can be used to implement this. I extend my support on adopting BGT CT.

Not approve CT: 1 WG member, 2 CAR co-authors

[Zafar Ali] Following the discussions during the IDR sessions and the list, the CT solution for multipathing appears to have basic control and forwarding design issues – if it needs to signal multiple routes for the same LSP. What is the purpose of flooding this duplicate information to all nodes ?]

[Dhananjaya Rao (DJ) CAR-co-author]: Opposes due to BGP-CT since uses the same encoding as BGP-LU (3107/8277). See additional comments in:

https://mailarchive.ietf.org/arch/msg/idr/R-ockSQHvDojiaEu_OgLBdOxBk/

[Swadesh Agrawal, (CAR Co-author)]: In addition to DJ's reasons, Swadesh lists 4 other reasons BGP-CT should not be approved. See his full text in:

<https://mailarchive.ietf.org/arch/msg/idr/AB9pSxxfG0fwOT0gBuiej9W8mpQ/>

Three Questions:

1. Do you agree or disagree that these two drafts are functionally identical?
 - [Aravind]: I agree [functionally the same], but I believe BGP-CT provides the functionality in much more holistic way.
 - [Reshma]: They look similar but there are minor differences. From customer engagements we learn that there are some use cases that BGP-CT solves that CAR does not address. This highlights there are some functional differences.

CAR and CT Adoption Poll Question 2 – Week 1 Summary (7/6 to 7/14)

- [Moshiko]: Both provides similar functionality, but BGP-CT is easier to implement, troubleshoot and educate because it use an existing procedures.
 - [Bruno, CAR co-author] I feel that those two drafts are mostly functionally identical so I think that we should rather have a single draft.
 - [Swadesh, CAR co-author]: Agree but CT is based on workarounds to even provide base multipath/protection which has implications.
 - [Israel, CT co-author]: CAR and CT are functionality similar therefore I support moving forward with a single draft.
 - [Gyan, CT co-author]: Agree
 - [Balaji, CT co-author]: I think the drafts attempt to solve the same problem, namely extending service mapping across domains.
2. If you agree, should we have just one draft or do the operational difference encourage us to have two drafts?

Single draft:

- [Aarvind]: I agree. We should converge towards a single draft
- [Reshma]: One Draft, and BGP-CT would be my pick.
- [Moshiko]: I think we need one draft.
- [Bruno, CAR co-author]: I feel that those two drafts are mostly functionally identical so I think that we should rather have a single draft.
- [Swadesh, CAR co-author]: Just need one draft. Support BGP CAR for protocol, scale and operational similarity to BGP LU.
- [Israel, CT co-author]: CAR and CT are functionality similar therefore I support moving forward with a single draft.
- [Gyan, CT Co-author]: Only one draft should be adopted for interoperability. CAR & CT are functionally identical, however adopting both drafts could result in interoperability issues.
- [Balaji, CT co-author] think having one draft will avoid duplication of effort and minimize interop issues.

Multiple Drafts

- Robert Razzuk:
Watching this debate I think the best option is to accept both proposals as experimental. Let the market decide not the mailing list then move one of them to Standards Track and the other one to Historic.
- Jim Uttaro (from Adoption Forum 1): I support the adoption of both BGP CAR and BGP CT.
“...Let the market decide not the mailing list then move one of them to Standards Track and the other one to Historic.” (Robert Raszuk)

I agree [with Robert]. Currently there are two solutions although they are “functionally identical” there are important reasons why operators may opt for one or the other. I am a bit confused as to the

driver “We need one interoperable solution”. Kompella and EVPN both can be used to create PWs, Multicast can be deployed using ingress replication, Rosen etc.... The time to drive towards one interoperable solution is before the technologies are specified. I look forward to a discussion of the use cases addressed in order to better inform architect/designers as to which approach is most appropriate for them. Letting the market decide will provide needed perspective to resolve this in favor of one or both specifications.

3. If you disagree, do the functional differences encourage us to have one or two drafts adopted?

[Swadesh, CAR co-author]: Just need one draft. Support BGP CAR [due to CAR’s] protocol, scale and operational similarity to BGP LU.

[Gyan, CT co-authors]: CT & CAR provide similar functionally, however the mechanisms and inner workings are vastly different. Therefore, only one proposal should be adopted to ensure interoperability.

- Both CAR & CT act at both the Service layer & VPN layer and provide mapping of VPN service route color to transport layer technology using next hop resolution schemes.
- CT provides an overall broader sweep of technologies supporting RSVP-TE and all inter-as options widely deployed today as well as Segment Routing SR-MPLS & SRv6.
- [Gyan’s comparison and contrast]
<https://mailarchive.ietf.org/arch/msg/idr/S39ZABGn4Ny0khKhWkJbgcENmhs/>

[Balaji, CT co-author]: I’m inclined to having just one draft, and my preference would be BGP-CT. I prefer BGP-CT for the following reasons (see his reasons at):

- <https://mailarchive.ietf.org/arch/msg/idr/mZEyLrvb2ooXDvGMryznZBac7BE/>