

CAR-CT Adoption Poll – Part 2: Adoption Detailed Summary (7/6 to 7/25)

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/>)

Call for more use cases:

- **[Aijun Wang]:** It seems that we need more additional comparisons based on the same use cases and topology that described in <https://datatracker.ietf.org/doc/draft-hr-spring-intentaware-routing-using-color/> to make better judgment on the selection of different approaches.
 - As operator, we are glad to select one solution that can interoperate with other among the key vendors. I think this can also reduce the implementation efforts for the vendors to support both of them.
- Regarding to the current two approaches, the followings are my thoughts and concerns:
 1. The key difference for CT and CAR is where to put the color/transport class/intention within the BGP encoding. CT utilize the existing VPN infrastructure, CAR design one new NLRI encoding.
 2. Then CT can utilize some existing mechanism (RTC) to filter the unwanted colored route, but CAR need to invent some new tools to achieve the same goal.
 3. CAR can be easily accepted/incorporated with the SR/SRv6 Policy, whereas CT can easily be understood with the wide deployed RD/RT based VPN services.
- There are also some others technical differences between these two approaches, but as operator we have concerns more on the deployment scalability, operation complexity and the pressure on the edge, transit and border routers.
 - For example, if the transport class in CT can be triggered automatically based on the received service intent, it can certainly release the operator from tedious configuration.
 - Let's continue the comparison on the 3rd part of this adoption call that raised by Sue. (see https://mailarchive.ietf.org/arch/msg/idr/d5V1i2_bGeDxeMhm6kd_sppfXOA/)

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.

[Miya Kohno]: I support the adoption of BGP-CAR.

[Phil Bedard]: I support adoption of the BGP CAR draft.

- From a former operator perspective, I would rather see a modern approach to solving the use case. I do not believe re-using technology from existing L3VPN makes BGP-CT easier to operate or deploy, it adds unnecessary complexity and encumbers it with the existing technical issues of L3VPN. A solution following more closely to the existing seamless MPLS solution using BGP-LU makes sense.
- The base constructs should not be tie to a specific data plane like MPLS, any encoding/encapsulation should natively support multiple data plane types.
- SR/SRv6 are seeing widespread adoption with the color community defining intent. This same color community should be carried through to a multi-domain solution. Having an additional field and mapping adds complexity and is prone to errors.

[Ahmed Abdelsalam]: I support the adoption of draft-dskc-bess-bgp-car. The draft provides a solution that offers extensible and very well-structured route format. In addition, it has interoperable multi-vendor implementations.

[Mankamana Mishra]: I support the adoption of draft-dskc-bess-bgp-car-05.txt and oppose the adoption of draft-kaliraj-idr-bgp-classful-transport-planes-17.txt. Though many points discussed and documented at https://mailarchive.ietf.org/arch/msg/idr/4Sffm_AR-wW9_jEEn3UfPGBbG4U/, some of the reasons for my input are:

- As WG we must have solution where future extension can be easy. In case of BGP CAR solution multiple transports including MPLS , SR-MPLS and SRv6 are supported natively. All of

CAR and CT Adoption Poll Question 2: Summary for Weeks 1-3 (7/6 to 7/25)

these transport are very well deployed or in process of deployment. Native support provides flexibility to be able to deploy different transport without much impact where as CT solution supports only label encoding in the route NLRI. Addition of any other encap is not handled automatically. This point does make BGP CAR as robust solution compare to CT.

- Operational simplicity : BGP CAR being similar to BGP-LU in its operational model, its preserving the BGP-LU simplicity .
- I am also aware of BGP CAR implementation by multiple vendors including cisco .

[Julian Klaiber]: In my master's studies at the Eastern Switzerland University of Applied Sciences, I'm working heavily with Segment Routing and also MPLS. I think the coexistence of these two protocols is very important, especially when moving from MPLS to SR. Any BGP transport solution must recognize this and also support multiple types of transport natively. Efficiency and very importantly also scale should not be impacted. BGP-CAR is a solution to this problem.

[Carmine Scarpitta]: a PhD Student at the University of Rome "Tor Vergata". I'm working on the implementation of BGP SRv6 L3VPN in FRR and SONiC. I support the adoption of draft-dskc-bess-bgp-car. BGP introduced a very simple and efficient operational model for multi-domain and multi-AS routing. BGP CAR adheres to the same routing model and proposes a mechanism to signal the intent-aware paths across a multi-domain transport network.

[Severin Dellsperger] Currently, I'm working as a research engineer in the field of Segment Routing and appropriate technologies. I think many networks will run SR-MPLS and SRv6 in parallel. The BGP-CAR technology helps to support multiple modes of transport natively and efficiently. Furthermore, BGP-CAR is totally compatible and consistent with the SR-Policy solution which is well-known. So, I support the adoption of draft-dskc-bess-bgp-car.

[Dhamija, Amit (Rakuten): I've read the documents and from Rakuten side we support the adoption of the CAR draft.

[Daniel Bernier] (Bell Canada): As an operator deploying SRv6, and already using color-aware SRTE, I welcome the effort towards standardization and consistency of the policy architecture. As such, I support adoption of the BGP CAR solution for the following reasons.

[Daniel Voyer] (Bell Canada): As an operator who has deployed SR SRTE, I find BGP-CAR to be consistent with SR-Policy architecture, and its use of the Color based automated steering a better bit for our network deployment. I support adoption of the BGP CAR solution.

[Stephane Litowski] I support the adoption of BGP CAR. Speaking with with my long experience of large network operations and engineering, I find BGP CAR better because:

- it's more straightforward in term of operations: with BGP services on top (which is the primary/only case), it becomes really similar to BGP services over BGP LU that seamless MPLS networks are using today. It is also allowing to reuse the same "autosteering" mechanism that SR-TE uses today (BGP service route with color ext community C and NH E is steered over a transport path that satisfies C and has an endpoint of E)
- I don't see the point of introducing VPN like mechanism for transport. It brings more confusion to me than helping the operations.

[Note from Chair: This is a technical opinion from Stephane]

- from a technical point of view:
 - As the NLRI uses the same components as SR-TE policies, it integrates very well with SR networks while at the same time still working with RSVP-TE or any other

- technology. But the trend of the industry is towards Segment Routing so this is good to design in such a way while keeping compatibility with older technologies.
- BGP CAR learns lessons/mistakes from the past and provide a more flexible encoding (supporting various dataplanes) while keeping packing efficiency. While designing as legacy and repeat mistakes when we can do better ?
 - I hear the arguments on the list about what's happening when there are network interconnections with different administrations of colors. BGP CAR addresses it through the use of LCM. While nobody can say that this scenario cannot happen, people cannot say that this will be the norm. If the two networks are under a completely different administration point of view, likely there will not be a BGP CAR interco but rather a BGP service interco like option A because of security/non-trust issues (even option B is sometimes challenging because of security purposes). If networks are interconnecting using BGP CAR, likely they trust each other and likely belonging to the same company and there are a lot of cases where color usage can be synchronized (company-wide network design guidelines.). Of course, there are cases, where it's not possible (independent affiliates.) and this is where LCM/color mapping can be used, but this makes it a corner case, not the norm. So, it's good to have a very simple solution that doesn't use any BGP community and that will address most of the cases, and "corner cases" are addressed using LCM.

[Stefano Salsano] (researcher, University of Rome (?)): I support the adoption of draft-dskc-bess-bgp-car-05.txt. I have been working on SRv6 standardization and implementation, the BGP CAR solution nicely addresses the *interoperability with multiple transports* including SRv6.

[Antonio Cianfrani (researcher, University of Rome (?)): I support the adoption of BGP CAR draft (draft-dskc-bess-bgp-car-05) since it is based on the same routing model of BGP for IP routing, enhancing it with color extension, and it works seamlessly across traditional networks. Moreover, SRv6 is used (or it is planned to be used) by many network operators and CAR provides BGP to efficiently support multiple transport solutions.

[Laurence Metzger] (OST.ch): BGP for IP routing and BGP labeled-unicast for seamless MPLS have set a very simple and efficient operational model for multi-domain, multi-AS routing. BGP CAR adheres to the same routing model and extends it with color.

[Rabadan, Jorge] (Nokia): I support draft-dskc-bess-bgp-car for WG adoption. From an NLRI encoding perspective, BGP CAR is a natural evolution of BGP LU, and allows for future extensibility. I believe BGP CAR is on the right track and it should be adopted by the WG. I also think the VPN encoding and use of RDs and route-targets proposed by BGP CT is not the right fit for a BGP transport AFI/SAFI.

[Gaurav Dawra]: (LinkedIn) Support - draft-dskc-bess-bgp-car for WG adoption.

[Wim Henderickx] (Nokia): I vote to have 1 draft and have the WG decide. If the market decides we will all implement them all at some point and I believe the IETF WG have a role to play to drive the industry in a certain direction. I support to adoption of BGP-CAR, since VPN machinery in transport is overkill.

[Reza Rokui]: (Ciena): I will support the adoption of following draft: draft-dskc-bess-bgp-car-05.txt. IMO although both drafts address the same problem, it would be better to have only a single draft.

[Ianik Semco]: (Cisco): I strongly support the adoption of the BGP-CAR proposal.

CAR and CT Adoption Poll Question 2: Summary for Weeks 1-3 (7/6 to 7/25)

- I agree with the many arguments brought forward by Robert R. on this thread along with the network operators' views expressed by Luay J. here: <https://mailarchive.ietf.org/arch/msg/idr/PKkqHGQYlrU2or3Af3ABbODq24I/> and by Daniel B. here: <https://mailarchive.ietf.org/arch/msg/idr/2tFGYFXsNtpPKSQdr8CCYFbt0tU/> .
- Also, aligning a BGP-based color-aware transport solution to BGP-LU seems an obvious choice (this is BGP-CAR). On the flip-side, bringing VPN awareness to each transport 'hop' can quickly become quite cumbersome (this is BGP-CT).
- <personal comment> I strongly believe the IETF should standardize on modern, scalable, and forward-looking solutions rather than stitched-together proposals that may be (granted) easy to demonstrate in a POC but are likely leading to operational complexity in highly scaled production networks. </>

[Pablo Camarillo] (Cisco): I support the adoption of BGP-CAR.

[Satya Mohanty (Cisco)]: Although functional differences are small, it is not clear to me how implementations based on respective drafts will inter-operate. Therefore, I will go with the CAR.

[Jose Liste]: I support WG adoption of draft-dskc-bess-bgp-car.

- As the WG embarks on the standardization of a new transport SAFI, it is crucial to start with the right base solution. This includes applying the lessons learnt from more than 20 years of deployment of its predecessor (BGP-LU) and avoiding shortcuts. In my opinion, BGP-CAR brings the right solution to elevate BGP to provide intent-aware inter-domain routing. Three key architectural advantages of BGP-CAR include:
 - BGP-CAR provides a natural evolution of BGP-LU that shares the same operational and routing models of Seamless MPLS networks. BGP-CT's use of VPN semantics for transport prefixes is overkill and unnecessary.
 - BGP-CAR provides automated steering and resolution based on color ext-communities applied to both service and transport routes. BGP-CT unnecessarily introduces different mechanisms like mapping communities that create operational overhead.
 - BGP-CAR provides intent-aware multi-domain paths across one or more encapsulation types (e.g.; MPLS and SRv6). Consider multi-domain networks with domains of different encapsulation or even domains supporting multiple encapsulations. BGP-CT proposal does not provide clarity on how to achieve this and break away from an implicit assumption of MPLS encapsulation.

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-author] 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.

[Jim Guichard, CAR co-author]: As a co-author I support adoption of draft-dskc-bess-bgp-car. BGP CAR as far as I can tell is consistent with existing SR-Policy solutions with direct leverage of Color Extended community for automated steering. CT, on the other hand, requires additional mapping. That leads to unneeded complexity.

[Yuanchao Su, CAR co-author]: As co-author, I support the adoption of BGP-CAR. We've been deploying <E,C> SR Policy model, and is extending this model to support per prefix multiple colors(intent), so a mechanism which is totally compliant with use of the Color ExtComm for automated steering of services with SR Policy is a must, as we don't expect a folk-lift solution.

[Ketan Talaulikar, CAR co-author]: [As a Co-Author], TL;DR - I support the adoption of the BGP-CAR draft. Arrcus has an implementation of BGP CAR and has done interop testing with Cisco. As the WG evaluates the two BGP proposals, it is essential to step back and review [draft-hr-spring-intentaware-routing-using-color] which captures the broader problem statement.

[\(https://datatracker.ietf.org/doc/draft-hr-spring-intentaware-routing-using-color/\)](https://datatracker.ietf.org/doc/draft-hr-spring-intentaware-routing-using-color/)

- What we need from BGP is to signal these "intent" aware paths for destinations through multi-domain networks. Color is an existing notion in BGP as an abstraction for "intent". Therefore, the most natural data model representation in a BGP NLRI would be (Prefix, Color). BGP CAR proposal is aligned with it.
- Multi-domain networks today use different encapsulations in different domains based on various considerations - MPLS, VxLAN, SRv6, GRE, etc. The "intent" aware paths can be set up to destinations across one or more combinations of these encapsulations. We need the ability to not only stitch these paths at borders but also simultaneously signal multiple encapsulations since some domains may support more than one (either during the migration or in a steady state). BGP CAR proposal is trying to tackle this in a clean way without any implicit assumptions of MPLS.
- This is a new problem space that we are trying to tackle. To that end, a protocol needs to evolve and develop for handling the new requirements.
- That said, BGP CAR is not the first AFI/SAFI that is introducing or dealing with an extensible NLRI design (e.g., EVPN). Many such mechanisms are widely implemented and successfully deployed. BGP CAR is a fresh/clean approach that leverages and learns from existing BGP mechanisms as appropriate.
- I am sure the technical review in the WG will improve the specification (e.g., the inputs on error handling) and make them more robust for not just CAR but other similar BGP extensions required in the future. Therefore, some of these opinions being raised on the NLRI and encoding design seem like red herrings to me. I am all for leveraging existing BGP mechanisms but trying to force-fit those mechanisms to solve an entirely different problem space or resisting clean design approaches can lead to the ossification of the protocol.

CAR and CT Adoption Poll Question 2: Summary for Weeks 1-3 (7/6 to 7/25)

- The BGP CAR keeps the notion of Color (as the "intent") consistent between the BGP Services and the BGP Transport layer. It maps in a straightforward manner to SR Policy mapping but is applicable to other "intent-aware" technologies like IGP FlexAlgo, RSVP-TE, and "best-effort" mechanisms like IGP-SR and LDP. There is a good balance of simplicity (I would call it "color consistency" across layers) and flexibility with policy mechanisms like mapping between colors, between layers, fallback mechanisms in resolution, etc.
- Most importantly, BGP CAR is exactly similar in operations to the BGP-LU (Seamless MPLS) design which most operators are familiar with. The BGP CAR proposal, only adds two things to the BGP-LU design - (a) the color ("intent") and (b) support for multiple encapsulations. All other aspects like route propagation, policies, convergence, scaling, and resolution are similar.

[Haibo Wang, co-author]: As a co-author, I support BGP-CAR solution. BGP-CAR uses (E, C) to describe an E2E tunnel with intent. This is the same as the idea of SR-Policy. Currently, all service routes use (N,C) for recursive tunnel resolution. Therefore, BGP-CAR with (E, C) as the key of the tunnel can match the intent of service well. For services that have evolved to the service with intent, the tunnel can be used more smoothly.

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/

[Krzysztof Szarkowicz]: I [support] the adoption of BGP-CT and the rejection of BGP-CAR.

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.

CAR and CT Adoption Poll Question 2: Summary for Weeks 1-3 (7/6 to 7/25)

- 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.

[Shraddha Hedge]: I support the adoption of BGP-CT as it uses constructs that are well understood by operators for many years and operationally simple and easy.

[Srihari Sangli]: I support BGP CT adoption. It leverages existing protocol mechanisms, operational aspects, easy to comprehend and troubleshoot in multi-domain networks.

[Julian Lucek]: I support adoption of BGP-CT. I have found that it is very intuitive to explain and operate, which is appreciated by the user-base.

[Jeffrey Zhang]: [I support the adoption of bgp-ct.

[Krzysztof Szarkowicz]: I [support] the adoption of BGP-CT and the rejection of BGP-CAR.

- BGP-CT introduces the concept of transports classes, which allow grouping of tunnels (not only BGP-CT, but all sort tunnels, like RSVP-TE, SR-TE, SRv6-TE, FA, ...) with similar characteristics (e.g., latency optimized tunnels) into common transport class, represented by a transport RIB. This simplifies considerably the operations, IMHO. Resolution policies allow to

CAR and CT Adoption Poll Question 2: Summary for Weeks 1-3 (7/6 to 7/25)

define flexible failover mechanisms, where the operator has full flexibility to define, what should happen if a tunnel with specific color is not available.

- Further BGP-CT leverages on well proven, widely deployed NLRI encoding, extended with additional attributes to carry the color. In general, BGP-CT is easy to understand, and easy to operate, and smoothly integrates into existing deployments.

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/2jkdJD4AP1pyOl7FwslA3KM8aZo/>

[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/0ZOHIDVii0MhlouedA0Q7EY0_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/S39ZABGn4Ny0khKhWkJbgcENmhs/>)

[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: 6 WG member, 3 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/AB9pSxxfG0fwOTOgBuiej9W8mpQ/>

[Ketan Talaulikar, CAR Co-author] I do not support the adoption of the BGP CT mechanism. It is trying to retrofit MPLS L3VPN constructs from the service layer to the transport layer and in doing so magnifies its well-known challenges (see [A] and [B]) unnecessarily. It obfuscates the notion of "intent" or "color" with the introduction of new terminologies and a requirement to always perform policy mapping at each router. The transport needs to be simple to manage and operate. I am cautious not to get lulled into the notion of familiarity with VPNs as this might be just the tip of the iceberg. I note that some of the authors of BGP CT have two other proposals ("associated" drafts as indicated by Sue) to try and provide a scaling solution with multiple MPLS label spaces (a new AFI/SAFI) and a new BGP attribute for multi-NH. So, my concern is that taking an IMHO "time to market" (?) approach for a very important problem space can cost us more in the long run.

[Note: Chair asked Ketan to take his concerns to the Forum 3 for an in-depth discussion.]

[Stephane (slitows)]: I support the adoption of BGP CAR. The two proposals are too close from a functional perspective to justify two solutions that would create more burden/complexity in the industry/deployments. Then I have two oppose to WG adoption of BGP CT

Antonio Cianfrani (researcher, University of Rome (?)): I don't support the adoption of BGP Classful Transport Planes draft (draft-kaliraj-idr-bgp-classful-transport-planes-17) since it solves similar use cases with lower efficiency than BGP CAR solution, due to its VPN-based mechanism. I think that there will be no gain in adopting another solution that basically supports the same functionality, so I support the adoption of one draft only.

[Mankamana Mishra]: I support the adoption of draft-dskc-bess-bgp-car-05.txt and oppose the adoption of draft-kaliraj-idr-bgp-classful-transport-planes-17.txt.

[Ahmed Abdelsalam]: I think there is no need to adopt CT draft.

[Pablo Camarillo] (Cisco): Also, I do not support the adoption of BGP-CT.

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.
 - **[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.
- **[Miya Kohno]:** Basically yes. But I would say they have different origins. BGP-CT is MPLS-native, whereas BGP-CAR is SR-native. From the SR/SRv6 operation point of view, CAR is simpler, more scalable/extensible and more consistent with SR Policy. BGP-CT's indirection and the need for the Mapping Community seems unnecessarily complex.
 - **[Robert]:** How about SR-MPLS ? Do you classify it under MPLS or under SR :) ? Don't you think CAR would be a good fit as well even for MPLS customers especially those LDP free and/or RSVP-TE free ? I would rather classify the two proposals a bit differently:
 - BGP-CT - of legacy data and control plane origin
 - BGP-CAR - modern data plane + extensibility
 - The choice is simple ...
 - *[Note: The Chair denoted that there was controversy here and Robert should move his comments to Adoption call Part 3.]*
 - **[Miya Kohno]:** Completely agree (with Robert).
 - **[Gyan Mishra]:** (to Robert) I would not call using L3 IP VPN semantics legacy as it is a proven technology that works very well has been used in the past for VPN overlay and will continue to be used in the distant future with SR and beyond technologies. L3 IP VPN technologies are far from being legacy or deprecated. We are seeing the technology being used in other areas such as
 - SD WAN and now at the transport layer due to its proven track record for operators.
 - BGP-CT control plane applies to existing technologies as well as all future technologies.
 - *[Note: The Chair denoted that there was controversy with Gyan and Robert should move his comments to Adoption call Part 3.]*
- **[Aijun Wang]:** OK. The more concrete description should be:
 - CT add "RD" field to the RFC8277 encoding (BGP-LU) to make the encoding similar to the existing VPN prefixes encoding, while CAR put the color directly into the NLRI key. In my opinion, the service intent should be decoupled from the underlying transport class division. As discussed in the list, the service intent may be described as some combinations/preferences of the transport class.
 - Should we slice the network in advance or on demand, and then steer the customer's intent traffic on them?
 - **Kaliraj [response to Aijun]:** confirmation of SAFI and VPN addresses]
 - https://mailarchive.ietf.org/arch/msg/idr/GYXaStRs2xVs5r6MT9bFszU5E_g/
- **[Shraddha Hedge]:** Yes agree the two drafts are functionally identical

- **[Jim Guichard]:** Yes, I agree with this assessment.
- **[Phil Bedard]:** Yes, these drafts deliver the same high level function
- **[Shraddha Hedge]:** yes, agree the two drafts are functionally identical
- **[Ketan Talaulikar]:** I believe they are functionally identical for the most part. However, the support for multiple encapsulation signaling is not very clear or obvious (at least to me) in the BGP CT draft. Other similar extensibility mechanisms may also arise down the line.
- **[Stefano Salsano]:** (SRv6 Researcher) I think that there will be no gain in adopting another solution that basically supports the same functionality, so I support the adoption of one draft only.
- **[Laurent Metzger]:** Yes
- **[Mankamana Mishra]:** Agree they are functionally identical. But there are differences such as protocol encoding and compatibility with SR-TE
- **[Rabadan, Jorge]:** I agree
- **[Gaurav Dawra]** (Linkedin): I agree
- **Wim Henderickx** (Nokia): I agree
- **[Reza Rokui] (Ciena):** Both drafts address the same problem
- **[Ianik Semco (Cisco):** Yes, I agree with this assessment. However (see my comments above)
- **[Stephane Litowski]** I agree that they are functionally identical in term of goal that they achieve (intent/color based routing) but in term of protocol encoding/efficiency they are different and BGP CAR has a better design. I don't know if you account protocol aspects in the functional part or not (functional part of the protocol).
- **[Ahmed Abdelsalam]:** Even though the functionally identical. BGP CAR offers more extensible and very well-structured route format.
- **[Pablo Camarillo] (Cisco):** I agree they are functionally identical, but their design is completely different. As already argued by others, BGP-CT seems to be a more narrowed solution from extensibility point of view. Also, I'm not at ease using the VPN encoding for a BGP transport as proposed by BGP-CT. On the contrary, BGP CAR is a natural evolution of BGP LU with future extensibility. Hence my preference for BGP-CAR.
- **[Jeffrey Zhang (Juniper)]:** They're largely the same, though there are some additional functionalities that bgp-ct provides, as pointed out by Kaliraj/Natrajan.
- **[Satya Mohanty (cisco)]:** Yes, the two drafts are functionally very identical. They address the same problems.
- **[Jose Liste]:** I reviewed Jeff H. note and believe it to be accurate. Broadly speaking, the proposals are functionally identical.
- **[Krzysztof Szarkowicz]:** Yes The[y] are in practice providing similar functionality.

Disagree:

- **[Amit Dhamija]:** Disagree as CT is defined with limited scope.. However CAR is evolution of BGP-LU solution and extended to the SR-MPLS, SRv6 and compatible with SR-policy, IGP-FA.
- **[Julian Klaiber] (researcher):** The two solutions cover the same use-cases, but there are protocol and important operational differences such as the ones I mentioned in my previous email.
- **[Satoru Matsushima]:** Many of people already mentioned with different words, but let me provide my comments on the questions as an operator:
 - 1. Disagree on that these two drafts are identical since color exists in the NLRI key part or not is fundamentally different for basic path selection.
 - 2. n/a
 - 3. One draft should be adopted with existing policy enhancement to keep the same path selection key, and thereby CAR should be adopted.
- **[Ahmed Bashandy]:** Disagree – Design choices are different which matter[s].

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. 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/>
- **[Miya Kohno]:** One draft is preferable. And BGP-CAR is a more natural extension, especially for SR/SRv6, which is characterized by simplicity.
[Shraddha Hedge]: We should have one draft so that its easier for vendors to implement one solution and operators to understand. And deploy one solution. Otherwise the cost that need to be borne by both vendors and operators is huge.
- **[Phil Bedard]:** Yes, these drafts deliver the same high level function. We ideally should have one draft
- **[Shraddha Hedge]:** We should have one draft so that its easier for vendors to implement one solution and operators to understand. And deploy one solution. Otherwise the cost that need to be borne by both vendors and operators is huge.
- **[Ketan Talaulikar]:** One draft is definitely preferable. Don't disagree but will await clarification on the extensibility and encapsulation-related functionality.
- **[Amit Dhamija]:** One Draft that addresses all the requirements. One draft adoption to cover all the current & new requirements.
- **[Julian Klaiber]:** I think the BGP CAR is the better solution to adopt and there is no need to adopt CT.
- **Antonio Cianfrani** (researcher, University of Rome (?)): I think that there will be no gain in adopting another solution that basically supports the same functionality, so I support the adoption of one draft only.
- **[Satoru Matsushima]:** One draft should be adopted with existing policy enhancement to keep the same path selection key, and thereby CAR should be adopted.
- **[Laurent Metzger]:** I support the adoption of draft-dskc-bess-bgp-car (one draft).
- **[Mankamana Mishra]:** One draft would be best option to go with. And as mentioned earlier the one draft which should be adopted is BGP CAR.
- **[Jorge, Rabadan]:** One draft. It would simplify things.

- **[Gaurav Dawra] (LinkedIn):** BGP CAR draft as that would simplify adoption and deployment.
- **[Wim Henderickx] (Nokia):** I vote to have 1 draft and have the WG decide. If the market decides we will all implement them all at some point and I believe the IETF WG have a role to play to drive the industry in a certain direction. I support to adoption of BGP-CAR, since VPN machinery in transport is overkill.
- **[Reza Rokui] (Ciena)** IMO although both drafts address the same problem, it would be better to have only a single draft.
- **[Ianik Semco (Cisco)]:** We should pursue a single draft, BGP-CAR.
- **[Ahmed Bashandy]:** We should have just one draft (BGP-CAR)
- **[Stephane Litowski]:** We should have a single draft which should be BGP CAR
- **[Ahmed Abdelsalam]:** I would encourage adopting only one draft (the BGP CAR draft).
- **[Pablo Camarillo]:** One draft (BGP CAR)
- **[Jeffrey Zhang]:** We should have just one draft, and that should be bgp-ct. It seems that while functional differences are small, there have been two camps distinctively formed around the two solutions.
- **[Swadesh, CAR co-author]:** Just need one draft. Support BGP CAR [due to CAR's] protocol, scale and operational similarity to BGP LU.
- **[Satya Mohanty]** I believe it is better to have one draft. Below are some of my observations.
 - BGP CAR model is very similar on the widely deployed BGP LU model of “underlay” connectivity. This is opposed to CT that imposes VPN like semantics in the under-lay. CT has an implementation and provisioning issue that the same intent (ECMP) but from two unique origin-points (anycast/multi-homing) need unique NLRIs (two different Route Distinguishers). They need to be brought together at an ABR and then advertised upstream. Question is which RD to choose or whether to assign a new RD? Bringing routes together (via import) operation involves some work and therefore designating a new RD with possible new net and path (although this is really implementation dependent). In CAR this all already accounted for as it follows the LU paradigm, same net resulting in two different paths.
 - In CAR Convergence is addressed naturally as both paths can be programmed in the forwarding as ECMP/PIC as may be the case. If both next-hops need to be communicated upstream, add-path can easily be applicable.
 - Semantics of clean separation of contiguous key elements and non-key elements with explicitly specified lengths makes CAR very extensible and requires no RR upgrades. IMO, this addresses a historical fault inherent in typed-NLRI SAFIs like EVPN wherein with any new route-type addition it is a pain to upgrade the RRs.
 - Recursive-Resolution can easily be done with one color-aware route resolving over another color-aware route and other in fact other routes viz. learnt via SR policy, IGP Flex-algo, BGP-LU etc. [I realize that CT also has the desired handling here].
 - There are other points already discussed in the mailing lists by others.
- **[Jose Liste]:** I am of the opinion that both customers and networking vendors would benefit from having a single solution to address a given problem/use case. And my support goes for BGP-CAR
- **[Krzysztof Szarkowicz]:** Yes. We should have just one draft and it should be BGP-CT (not BGP-CAR)

Multiple Drafts

- **[Robert Raszuk]:** 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.

- **[Jim Guichard]:** This is the key question. While both solutions are functionally equivalent, operational differences are important from a deployment perspective. It is my belief that BGP CAR fits the deployment models I have seen (hence I am a co-author) and from a complexity standpoint is a better fit in those environments. Having said this there have been several examples in the past of having two solutions to solve the same problem although that is not my personal preference and I would therefore prefer to see only one draft adopted.
- **[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/>

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