

Section 2: Detailed Summary CAR-CT Adoption Poll 1 [Chair's notes]

General comments:

Response to Gyan's definition.

[Ketan]: The Gold, Bronze and Silver classes of the past may be somewhat limited and may not cover a wider set of “intents” that are now required [to be] address[ed]. When referring to QoS, perhaps PHB is what is being referred to? Which is again limited. Not all intents are related to QoS.

[Ketan]: The notion of color was introduced in BGP by RFC5512 and predates SR. That said, I do agree that “color” based steering of BGP services over SR Policies using Color Ext-Comm is probably the most widely implemented and deployed solution that leveraged “color”.

[Luay]: These are congestion control mechanisms and not sufficient for today's needs. We will always have QoS but it's one piece of the pie since many of the steering use cases needed are not QoS dependent in an IETF-based network

Response on Sue's view on Intent: There is no clear and universal intent-based language in 2022.

[Kaliraj:] BGP implements the Internet's policy and economic infrastructure. Enabling customers and providers to express SLA in BGP allows for Automation and further facilitates such economic interactions. It allows [networks] to create a slightly better feedback loop between inter-domain networks and avoids over provisioning of provider networks. 5G topology is one use case.

[Ketan]: I agree [with Sue on intent] and this is very important. We need to be careful when using the terms “intent-aware” or “intent-based” in IDR documents.

[Bruno]: +1 (on no clear language for intent.)

[Aravind]: There are several applications I have heard from my customers I work with where intent (color) is expressed to treat traffic flows in a particular way. For example, in order to service chain or to offer security services we would need to redirect them to the security component (firewall). In order to [handle] this at a granular level we can leverage BGP flowspec (as Kaliraj points out in 2c on ingress). The path to service chain component itself can be on any transport (RSVP-TE/SR-TE) and the SLA treatment is given to a particular flow. So I (and my customers) think Flowspec redirecting traffic over transport-class aware paths is a useful and desired feature.

Natrajan (NV)> I agree that intent is more than a color. Seamless MPLS networks have been able to guarantee intent using SRLG/admin-groups within a single TE domain. In Brownfield RSVP-TE deployment, there are adequate ways to do the above and I see Greenfield SPRING networks catching up to the same. However, BGP did not produce a holistic way of cross cross-connecting these TE characteristics that are necessary for guaranteeing the SLA end to end until now. BGP needs an identifier to inter marry these TE characteristics across multiple TE domains. That identifier is the color.

Natrajan (NV)> The second part is how flexible it [Color] is for services to express intent (language)?

[Sue]: IRTF document on definitions of Intent
(<https://datatracker.ietf.org/doc/draft-irtf-nmrg-ibn-concepts-definitions/>)

So let us use color as BGP defined value and transport class as BGP defined value.

[Ketan]: Correct. So, color is an existing BGP-specific abstraction of an “intent” (English word). [BGP-specific color exists in] ... SR Technologies (implemented, deployed, and understood [by those deploying it]). However, there is nothing that precludes “color” based steering of BGP services over other tunneling or encapsulation technologies.

[Bruno]: +1 (plus BGP already has a color extended community, and BGP-SR policy uses color so I think there is not a need to define a different term.

[SA] Yes. Color has been used to both express the desire to go over certain tunnel by attaching color extended community to service route (colored route) and to create paths/tunnels that satisfy requested intent called color aware route (example SR policy architecture).

Q1: What is the customer need driving the use of Color to express Customer Intent?

General responses:

[Halpern] (indirect) Want better mechanism to set-up paths across AS for a single Intent/QoS.

[Moses]: I believe the primary motivation here is the need for the ability to program the network (Network Programmability). Service providers need the ability to program different network paths that are optimized for Service Level Objectives(SLOs) and Service Level Expectations (SLEs) contracted by the customers.

- SLOs are directly measurable indicators such as guaranteed bandwidth, latency, packet loss rate etc.
- whereas SLEs are not directly measurable attributes such as path diversity/”disjointness”, security, encryption, isolation, geographic restrictions, etc.

Providers are interested in developing the capabilities to provide differential carriage treatment to achieve the SLOs and SLEs within a single domain, across multiple domains and across multiple autonomous systems. Tagging the service routes can be performed by the provider's service nodes customer at attachment points, customer end devices or even by a SDN controller that monitors the end-to-end performance of service/application.

[Moses]: Some of the use cases:

- Transport Network Slicing to support 5G Slicing (URLLC, eMMB, mMTC, etc.)
- Dual-plane/ Multiplane forwarding for isolation between primary and secondary traffic flows
- Dedicated network slice for high value wholesale customers
- Carriage of Edge Compute services with stringent SLAs

[Luay]: [this motivation] Could be driven by customer or by the infrastructure to treat traffic with specific SLOs. Network slicing is driving more of these use cases. However, it's more than QoS and most of your questions assume QoS is the main mechanism.

Question 1-a: Are applications requesting to be able to tag their routes with SLAs (color) at the service level?

[Ketan:] Yes. However, the tagging may be done by customers OR by operators for their customers/applications. This is what is widely implemented and deployed today with SR]

Week 1 – Summary of Car/CT Adoption Poll 1

[Bruno]: Different applications/customers have different requirements in terms of routing paths. [For example, they] may optimize for bandwidth, vs. link delay, etc. [One usual niche use case] is two different planes.

[Bruno] Note that we already have the need and solution within the single IGP (e.g. FlexAlgo, SR policies). IMHO the question is extending this to BGP, in a way allowing each AS to use its own solution for its Intra-AS Routing. Such AS RSVP-TE in one AS, FlexAlgo in another AS, [and] SR-policy in another AS.

[Natrajan (NV)]: Both a and b are true.

[Swadesh (SA)] Yes, routes are getting tagged to send traffic over path that satisfy some intent. This tagging can be done by operators or by their customers. Intents means more than just QOS. It could be low latency, avoidance of resource (link, node or domain) etc. Measurement may be based on intent type.

[Gyan]: Yes.

- Tagging has been done via [the] QOS model providing Value added service tiering for 20 years with Gold, Bronze, Silver SLA based strictly on traffic classification and PHB queuing mechanisms.
- SLA based on “coloring” meaning steering overlay to underlay slice underpinning had started initially with RSVP-TE steering.
- L3 VPN services with per VRF TE mapping next hop re-write which requires a separate loopback per SLA / Color.
- With the advent of SR-TE the concept of VPN service route coloring expanded to micro and macro flow “intent based” steering.
- Network slicing is a major use case that spans beyond just 5G to Enhanced VPN+ service offering which can use SR-TE / Flex Algo to color the underlay.

Question 1-b: If so, is it due to QoS/SLA measurements on traffic between Data Center applications and user applications (such as applications on phone)?

[Ketan]: We have requirements (“intents”) like low latency, low loss, best-effort, specific bandwidth guarantee, avoidance of certain parts of the network, etc. Not all “intents” need to manifest in QoS Policy or require measurements.

[Bruno]: This may either dynamic measurement or static attributes. Immediate use cases use static attributes.

[Natrajan (NV)]: Both a and b are true.

[Swadesh (SA)] ... Intents means more than just QOS. It could be low latency, avoidance of resource (link, node or domain) etc. Measurement may be based on intent type.

[Gyan]: Yes. What operators have found is that QOS traffic classification is not enough and being able to provide a way to take VPN service routes and somehow map them to different paths based on cSPF TE constrains such as delay, bandwidth etc is critical for voice and video delay sensitive and mission critical customer traffic.

Question 2: In the distant past QoS was hard to set-up seamlessly as a QoS pathway across multiple Autonomous systems (AS).

General comments:

Week 1 – Summary of Car/CT Adoption Poll 1

[Bruno]: We probably need to distinguish Multiple AS versus multiple Administrative Domains. A single Administrative Domain may cross multiple ASes, e.g. due to organization boundaries or technical scaling. I would hope that agreeing on a QoS/color within single administrative domain should be doable regardless of the number of ASes. Plus, in my view, a color does not match to packet forwarding QoS treatment ('a la diffserv) but rather to a to a [set of] different routing paths toward the same destination (NLRI).

[Luay (LJ)]>> Seamless could be accomplished in many ways, signaling is a big part of it. It's not a QoS pathway, I'd prefer to call it the SLO path that would have QoS treatments at different points. These questions are still QoS centric so nothing more to comment].

[Moses ([MN])] As I mentioned above, the intent is more than QoS. It can also be Service Level Expectations such as path diversity/"disjointness", security, encryption, isolation, geographic restrictions etc. E.g.: A media production company can stream multiple copies of the same live streams to the production site for resiliency. The intent can be signaled as stream A and Stream B (and Stream C) should be transported via mutually disjointed paths (within each AS's and/or across multiple AS's).

[Swadesh (SA)]: Yes it makes sense and already has deployment experience with SR policies.

[Luay (LJ)]: Seamless could be accomplished in many ways, signaling is a big part of it. It's not a QoS pathway, I'd prefer to call it the SLO path that would have QoS treatments at different points. These questions are still QoS centric so nothing more to comment.

Question 2a: Should Customer intent that expresses pathway "QoS" be passed in BGP routing updates sent between Autonomous Systems?

[Kaliraj: Yes]

[Ketan:] I'm not sure what is meant by QoS and "QoS pathway" here. Can you clarify what specific QoS attributes are being referred to?

[Ketan]: The point of setting up a color-aware transport path between points A and B in the network is to enable a BGP Service like L3VPN to operate between those two points (e.g. [so] A and B are seen [operate] BGP next-hops for the SAFI 128. This [design] puts [color-aware] "multiple-AS" [paths] into the right perspective – as in option C. Otherwise, Option A/B is used for services between ASes and CAR/CT reachability ends at the AS boundary.

[Bruno]: Yes IMHO for:

- the service routes (e.g. BGP or "internet" to express the requirement, and
- the transport routes fulfilling the requirement.

Note that the above already requires agreement on the meaning of Color between ASes and between administrative domains. (and many more than the/[this?] number of ASes may be involved to set-up interop domain transport colors.

[Natrajan (NV)]: yes

[Swadesh (SA)] Yes. [It] makes sense and already has deployment experience with SR policies. I don't see anything specific for flow specification. We should allow actions to redirect traffic onto these color-aware routes.

Week 1 – Summary of Car/CT Adoption Poll 1

[Gyan]: Yes. ... [And] this is possible today with SR-TE policy candidate path coloring encoded in BGP TEA not NLRI. I think this is an important aspect to be able to provide inter-as ... [for] end VPN service route coloring to underlay path being instantiated

Question 2b: Is it the purpose of color or transport class to allow automatic steering of traffic on into an “QoS” path (across different technologies)?

[Kaliraj]: Yes, See answers to question 1 [expression of SLAS, better automation, feedback loop].

- [Important] to cater to brownfield deployments with different technologies,
- [important] Preserve and extend ROI on time tested technologies and training.

[Ketan]: Color is an abstraction that is being passed in BGP updates. Whether or not it also indicates certain QoS parameters depends on the operator and their deployment design.

[Bruno]: Yes

Natrajan (NV): Yes. However, transport-classes provides a way to organize “QoS” path into separate buckets for easier service resolution,

[Gyan]: Yes. The goal of CT is to be provide a seamless mechanism of providing VPN overlay coloring to underlay loose underpinning to all steering technologies for both green and brownfield deployments.

Question 2c: How should this automatic steering interact with flow specification?

[Kaliraj]: Flow specification should be able to express SLAs just like other service families (IPv4 or LV3PN). For example, Color extended community carried along with Redirect-IP extended community.

[Ketan]: For this topic, we [IDR] should consider them [flow specification and color] orthogonal.

[Bruno] On my side, [the] initial use case does not involve BGP Flowspec. Automatic steering is expressed by the service/VPN routes attaching a color to express the color/path it wants to follow.

[Natrajan (NV)]: BGP features need to interoperate with each other seamlessly and allow [for] a deployment to leverage the benefits of the features being used. So, this is not only limited to BGP flowspec but to other BGP features as well.

[Swadesh (SA)]: ... I don't see anything specific for flow specification. We should allow actions to redirect traffic onto these color-aware routes.

[Gyan]: FS Redirect-IP Extended community plays an important role as a traffic classification tool that can be leveraged to map overlay VPN service route micro or macro flows to underlay constraint based paths.

Question 3: For those who believe that BGP should set-up a seamless path across multiple Autonomous Systems for a single Intent/QoS, do the exact mechanisms matter or do you simply want an interoperable solution? If they matter, please describe what matters.

Do not matter:

Week 1 – Summary of Car/CT Adoption Poll 1

[Halpern: Exact mechanism does not matter. I do want interoperability. Prefer one mechanism.]

[Bruno]: Primarily, we need interoperability. Then Automated steering should be simple and inline with what we already have with SR-policy and what we have within one AS. BGP encoding is less important, but for sure, it seems like different person have different opinions.

[Luay (LJ)]: I believe Bruno brought this up, we have different ASes under one administrative control on top of Different Ases of different administrative controls. Leaving QoS aside again, we want to be able to build an SLO path across the same AS, different ASes under the same administrative controls, and lastly ASes under different administrative controls.

Do matter:

[Kaliraj:] How, [BGP] carries this SLA information matters. SLA/Intent being an adjective should not be put in NLRI [since it is problematic] for TE,

[Ketan]: BGP signaling (and other mechanism) matters for interoperability. While some like QoS config and other implementation aspects like policies may not [matter].

[Bruno] Not having to manage unnecessary indirections or ID spaces would help IMO.

[Moshiko] Important things are:

- Easy to [experiment in lab with prior to deployment],
- Easy set-up with no/un-noticeable interruption to service traffic,
- Easy to troubleshoot and visualize
- Use of know mechanisms [e.g. RD and extended community] helps in traffic manipulation using regular policies, filters and roadmaps.

[Natrajan]: The summary of the IDR chairs is that BGP-CT and BGP-CAR are functionally similar. What matters most to the customer are the following:

- Do customers wish to carry color as a key in NLRI?
- Do customers wish to carry non-key “forwarding information” as part of NLRI?
- How would customers wish for their services to “Express Intent” for individual flows.

[Moses] Mechanism matters.

- Simplicity at the same time wide coverage of use cases - The solution framework must support all the current and future use cases in a simple manner. (It shouldn't be addressing only the typical use cases of subset of customers / 'lead operators').
- Mechanism should preferably leverage the familiarity of existing standards and mechanisms deployed by the customers.
- Mechanism should focus on the transport layer intent use cases rather than keeping the scope wide open for other non-transport use cases (e.g. allowing service layer intents to be mixed with transport layer intents)
- Mechanism should not allow and open a pathway to add any unknowns in future that could impact the stability of BGP.

[Sue]: Thank you for letting us know how key stability of BGP for routing protocols is to you. Your feedback seems to ask us to re-examine our decisions on using BGP for information passing as well as routing functions. Is this correct?

[Moses] My comment was about maintaining the functional separation and restricting the usage to the intended purpose rather than keeping it open.

- The purpose of BGP-CT/CAR AF is to exchange the intent-aware transport routes across domains. The mechanism shouldn't open access to the usage for other purposes.
- I do believe that BGP shouldn't be used for exchanging any unstable information.

[Swadesh (SA)]: Currently operators deploy BGP LU for best effort inter domain paths. Exact mechanism matter such that existing routing model, functionality and operational experience provided by BGP LU is maintained and enhanced.

[Gyan]: Yes. the mechanisms matter especially for OPEX and MTTR and network availability. Mechanisms used should allow for scalability as well as ease of use for MTTR and network availability. Use of TEA and path attributes is always better and more scalable when extended to CE as compared to NLRI encoding which can be problematic at CE layer adding too much complexity.

Question 4: RFC7606 focused on error handling in which the MP-NLRI focuses on destination keys (RD and Prefix) plus non-key material (Labels, SIDS). Attributes (generally) apply to all NLRI. For example, MED applies to all NLRIs in the packet.

Question 4a: Would error handling be better for color-aware routing if attributes relevant to a specific color/class be grouped in a MP-Color-Attribute?

Yes

[Kaliraj]: Yes. Not carrying “non-key fields” in NLRI helps in better error handling, ... [and] in keeping BGP modular [one AF, one modular function]. [Without] this focus on modularity [new AFIs tend to include previous AFIs] in to ““God’-AFI” ... with the subsumption of other functionality into ““God’-AFI.

[Bruno]: Probably, error handling would be better with a simple NLRI. However, this is a complex subject, and it really depends on the type of error and whether the length field match[es] the size of the value field (or not). However, there are other tradeoff[s] involve[d] (e.g. efficiency, extensibility/generality). Note that we have already encoded non-key data in the NLRI (e.g. label, label stack) so it may not be black and white.

[Natrajan (NV)]: Yes [to a and b] – BGP- BGP-CT team supports this approach as it confines the ... error handling to the attribute. Such an attribute is already being proposed by the BGP-CT team. I am bringing the same to the notice of both customers and the IDR chairs alike

<https://datatracker.ietf.org/doc/draft-kaliraj-idr-multinextthop-attribute/>

[Gyan]: Yes, that makes sense for BGP update packing efficiencies RFC 4277.

No:

[Ketan]: No

[Swadesh (SA)]: SA] Not sure if I see advantage of such attribute from error handling point of view. Each color aware route has associated attributes and hence color scoped. It provide error handling at granularity of each route.

Question 4b: Should IDR consider future work on a MP-Color Attribute?

Yes: (4b)

[Kaliraj]: Yes. [see] draft-kaliraj-idr-multinext-hop-attribute which can serve that purpose.

[Gyan]: As “coloring” VPN overlay to underlay colored path instantiation is critical for operators as they migrate to SR, having an MP-Color path attribute would be very helpful and could provide an SR-TE automated steering optimization and could help reduce SR Policy complexity as well as help with MTTR and availability.

No: (4b)

[Ketan]: I do not believe there is such a requirement for Color-Attribute. As a reminder, “color” is an abstraction.

No comment: (4b)

[Moses] No specific comments as I am not familiar with BGP error handling / RFC7606.

Bruno Discussion leading to a generic solution:

[Bruno]: Not sure what you have in mind. I don’t think that encoding the color is the issue as this is a short, constant size value. To me the difficult question is more about encoding some NLRI specific data (e.g. label for MPLS, but these days we may need to be more general) in order to avoid breaking packing (in which case, the question is not completely specific to BGP colors).

[Sue]: The question was on encoding some NLRI specific data (the second half of your comments). To [further explain this] consider:

- MP-REACH/MP-UNREACH - group NLRIs by AFI/SAFI
- An MP-Color attribute would group attributes by color and an An MP-NH attribute could group by NHs.

Example of an MP-Color attribute:

- [MP-Color-Green: {MED, Tunnels, non-key data (Next-Hop label-pairs, etc)}]
- [MP-color-Gold: (MED, Tunnels, non-key data (NH, Labels, backup, preferences)]

Example of an MP-NH attribute would group this by Next-Hops

[Next-hop-1: Color: Gold, tunnels, non-key data (labels, other things), preferences]]

[Next-hop-2: Color Green, etc.]

Please do not take my examples as anything other than vague directions IDR might go into.

[Bruno2]: Thanks for the clarification. The above proposal could help; however IMHO one difficulty is the selection of this additional "group by" key. You provided two (Next-Hop, Color), but there could be multiple. IMHO, the one which is likely to be the most useful (based on history and current discussion) is to be able to advertise non-key data per NLRI.

Week 1 – Summary of Car/CT Adoption Poll 1

[Bruno2] Personally, in order to be generic (and hopefully benefit from a single code base to debug) I would propose a new MP_REACH_NLRI attribute carrying sets of (NLRI, non-key data), compared to (NLRI) currently. Or may be even better for error handling two lists: a list of NRLI, followed by a list of non-key data (in the same order). But this may be seen as a larger effort compared to BGP-CAR which only defines a new SAFI.

[Jeff in Response to Bruno2]: This would nicely address the error handling point I'd raised on -car that Dhananjaya has responded to. The idea of a generic key that can safely make it through a route reflector without need to upgrade the reflector has strong appeal. We're seeing that it's challenging to do.

Also see Jeff Haas discussion on CAR encoding issue with Dhananjaya (DJ) Roa.

<https://mailarchive.ietf.org/arch/msg/idr/qorZr8Y7Cpn4cv-BkH3LhBL391I/>

(see text in full below with DJ's comments include]