

Computing-Aware Networking BoF (Non WG forming)

IETF-113, Vienna

Agenda

=====

Tuesday March 22, 2022

10am - 12pm (two hours)

Room: Grand Park Hall 1

Chairs:

Jeffrey Zhang zzhang@juniper.net

Linda Dunbar dunbar.ll@gmail.com

AD:

John Scudder jgs@juniper.net

=====

Administrivia - Chairs - 5 minutes

[John Scudder] do not discuss details today. More about what our next steps are.

[Linda Dunbar] please all speakers checking the notes to make sure everything is properly captured.

[Georgios Karagiannis] we had some side meetings, since IETF 109/110. There is already some discussion/debate on the use cases.

Problem statement, use cases and requirements

- MEC/CNC Update (Luis, 15 minutes) 20/120

[John] please provide the example of the scientific example.

[Luis] Transfer of datasets between research centers

[Acee] what is CNC?

[Luis] CNC stands for Compute and Networking Convergence (ITU-T-sg13). MEC is developed by ETSI. They are different initiatives and different architectures.

[George] what is your position in the edge location regarding to UPF?

[Dirk Trossen] (through chat) Hmm, apologies, realized you need TIES access for it (i.e. ITU account) - maybe there could be some open document the presenters could share on CNC?

[Linda] CNC/MEC is to show work by other SDO, e.g., ITU, not entirely for IETF

[Dirk Kutscher] (through chat) It would be useful to understand the relationship of these systems (MEC, CNC) with existing transport layer security and trust infrastructure for that. I think this was Erik's questions as well.

- Specific Use Cases (Peng Liu, 20 minutes) 40/120

[Rich Salz] (through chat) UE = user equipment?

[Dean Bogdanović] (through chat) yes

[Pete Resnick] (through chat) Perhaps a dumb question (I am from the upper layers and the air is thin and makes me stupid), but is there some sort of authorization model so that an edge can indicate whether or not it will provide compute services? I didn't see anything in the current drafts.

[Rich Salz] (through chat) It would appear to be the contract you signed with your phone provider.

[Tony Li] (through chat) I think what you really mean is service discovery and yes, that should be part of it.

[cabo] (through chat) (And the UE/App also needs to authorize ("trust") the edge to provide the service.)

[Dean Bogdanović] (through chat) anytime you talk to MNOs, they forget to mention SIMs and the AAA SIM provides

[Joel Halpern] (through chat) @Pete that is a good question. If one is trying to solve the useful question of how the application asks for services from the network. This seems (but I am not sure) to be something else.

[Marie-Jose Montpetit] (through chat) @ Pete: there is an edge discovery draft in COIN RG (it needs updating but some functionality is there)

[Cheng Li] (through chat) it will be a general requirement for MEC, even for services providing, but it worth to be discussed.

[Dirk Kutscher] (through chat) IMO, it's rather similar to today's edge data center deployments of OTT caches. The operator would decide to run such functions and then may route UE requests to those instances.

[Carlos Bernardos] (through chat) I wonder if mobility issues (and associated protocols) are also in scope. I think there are scenarios where routing alone would not be sufficient

[Mike McBride] (through chat) We've had had edge DATA discovery drafts in coin but not edge device discovery unless perhaps there is now both.

[Pete Resnick] (through chat) So service discovery will include some concept of permission (in both directions). OK, that makes sense.

[Tony Li] (through chat)@Carlos: Yes, absolutely. See Linda's draft.

[Joel Halpern] (through chat) @Carlos - if this were handled as an overlay question, I think mobility could be part of it. But trying to make underlay routing do any of this is a mistake.

[Dirk Trossen] (through chat) @Carlos indeed, routing in the light of mobility is clearly in scope - also looking at the use cases (wait for next slide)

[Carlos Bernardos] (through chat) @Joel: exactly, that's my point. I don't think routing can do that

[Mike McBride] (through chat) @Dirk. Sounds like this could be used as a high performance DLT network.

[Carlos Bernardos] (through chat) @Dirk: I'm not sure routing might be the solution for all the scenarios involving mobility. I think for some we might need routing+mobility solutions.

[Dirk Trossen] (through chat) @Carlos no problem, didn't mean any offense. My point is that if you want to ensure mobility on both ends (client and service), any routing decision (to a possible new service instance) must be complemented by ensuring the mobility of any transaction (dyncast calls this 'affinity') since routing to a new service instance may break ongoing transactions.

[Christian Hopps] (through chat) will there be a time to question the entire concept?

[Dirk Kutscher] (through chat) There is a discussion slot at the end of the session.

[Tony Li] There's open discussion later

[Luis Contreras] (through chat) @Carlos: regarding ALTO, certainly it can play a role on the overall idea of compute awareness, but differently from routing, so more in the line of exposing compute capabilities available in the network (see for example draft-contreras-alto-service-edge)

[Cheng Li] (through chat) as my understanding, it would be based on service, a service will be identified by an anycast address

[Tony Li] (through chat) This doesn't sound like a clarifying question...

[Carlos Bernardos] @Dirk: thanks for the clarification. I think this is an interesting problem.

@Luis: OK, I'll check the draft

[Aijun Wang] (through chat) ALTO can't let every router within the domain to select the best MEC server

[Zhenbin Li] (through chat) @Wim Regarding your question, please to refer to the following presentation. It doesn't have to be application-based. It can use the anycast IP address for the location of the application.

[Marco Liebsch] (through chat) regarding the previous comment: 3GPP solves the discussed problem up to UPF, more is needed for complete e2e traffic treatment over the N6 reference point. Just one pointer where this has been discussed: <https://datatracker.ietf.org/doc/ht...fattore-dmm-n6-cdpd-trafficsteering>

[Zhenbin Li] (through chat) The CAN dyncast work is to depend on the network device to steering traffic other than the UPF.

[Zongpeng Du] (through chat) 3GPP is also researching the distributed architecture now. Perhaps, the position of UPF in future will be much lower in the network than now.

[Christian Hopps] (through chat) So use Anycast to find the nearest alive load balancer, now it's not a single point of failure and it's multiple site aware.

[Pete Resnick] (through chat) Traditional DNS might be a problem, but DNS-SD might be adaptable for this. I don't know it well enough to say for sure.

[Dirk Trossen] (through chat) @Pete Peng is pointing out the dynamicity of possible relations, so any indirection step adds latency when wanting to make decisions at high frequencies (e.g., due to load changes, mobility).

[Daniel Bernier] (through chat) one elephant in the room I think is expecting that edge computing applications will rely solely on DNS or equivalent. Current development work show that it is becoming tighter with patterns like service meshes in cloud with application traffic routing logic (ref envoy sidecars or equivalent). So applications are getting smarter and smarter at optimizing traffic routing between client and server. However, as per Wim's comment, the application might know where its best edge is from a workload perspective but mobile UP (gNobe to UPF mapping) might not necessarily be there.

[David Oran] (through chat) Question: This seems to assume conventional non-distributed applications just running at the edge. what about modern frameworks like Sapphire? and Ray?

[Christian Hopps] (through chat) hmm i just lost all AV.. this happened yesterday to me too, it was my side

[Zongpeng Du] (through chat) Load Balancer in this Layer 3 is explored in dyncast. In dyncast, the network should recognize the service ID (anycast IP) of the packet. The service ID is the destination address of the IP packet, it kind of contains some service information.

[Pete Resnick] (through chat) @Dirk: Understood, but I think DNS-SD or something else like Bonjour might have the anycast (or at least multicast) and dynamic properties that you need. And they are (as I understand it) direct to the device, not indirect. It may not fit, but it should at least be examined to see if it is useful or informative.

[Cheng Li] (through chat) good idea

[Dirk Trossen] (through chat) @Pete I would say that this should be captured in a more detailed gap analysis - so yes, agree.

[Zhengyuan Diao] Why compute resource is different with other resources?

[Peng Liu] It's not like networks that are measurable by bandwidth and delay. It is more complicated.

[Mengxiao Chen] We see more use cases, live streaming, and short videos, hopefully can be integrated. Livestreaming might also benefit from network optimization they should also be considered.

[Peng Liu] We are happy to join efforts.

[Wim Henderickx] Do we say we should select UPF based on the application? Steering is done per user? or per application?

[Peng Liu] Application might be slower than a routing solution.

[Wim Henderickx] We can solve the problem on different levels, not to preclude. We need to specify the explicit level.

[Greg Mirsky] Is computing resource measurable?

[Peng Liu] It is, and how to use the measurement would be solution related.

- Gap Analysis and Requirement (Peng Liu, 20 minutes) 60/120

[Hao Li] I don't know if there is other candidate solutions. say ?? Shall we also investigate this?

[Peng Liu] This is doable and can be integrated in future version.

[Qiang Dai] For the load balance,

Potential solutions

- Solution using existing techniques: (Shraddha Hegde, 15 minutes) 75/120

Load Balancer

[Aijun Wang] (through chat) The drawback of load balance based solution has been analyzed

[David Oran] (through chat) For the state of the art on L4 LB, check <https://dl.acm.org/doi/10.1145/3123878.3132012>

[Mat Ford] (through chat) Application layer invisible to L7 too presumably if apps running over e.g. QUIC

[Christian Hopps] (through chat) Wouldn't a L7 load balancer be working hand in hand with the application itself, so deep packet inspection isn't needed.

[Adrian Farrel] (through chat) @DaveO The question for me is whether this needs to be solved at the transport layer or at a lower layer. Doesn't it depend on the "scope" of the transactions?

[David Oran] (through chat) @Adrian - this is a "cross layer" problem, for a different take on how to decompose this check out <https://conferences.sigcomm.org/acm-icn/2019/proceedings/icn19-16.pdf>

[Zongpeng Du] (through chat) agree, there are some tradeoffs here 11:10:34

Layer3 load balancer(a router) can obtain the network information easily, and more is closer to the user. The whole network can work as a distributed virtual Load Balancer.

[Christian Hopps] (through chat) it doesn't scale

[Lars Eggert] (through chat) i'm trying to understand if there is anything missing from current lbs that would prevent their use as-is? other than there is for market reasons no interop standard between different lbs?

[Joel Halpern] (through chat) @Chris - and injecting all of this into routing does scale? No.

[Daniel Bernier] (through chat) so UE anchoring to UPF needs to happen first

[Christian Hopps] (through chat) No that is my point

[cabo] (through chat) LBs are inside the application's security bubble.

[Christian Hopps] (through chat) @joel, this entire concept being presented here doesn't scale

[Tony Li] (through chat) As I understand it, session persistence across migration isn't covered

@Chris why not?

[Dirk Trossen] (through chat) @Lars I wonder how to coordinate the LBs if you mix and match vendors. Isn't this related to what CAN proposes, namely address the needed signaling of metric information?

[John Scudder] (through chat) @tli It wasn't clear to me how hard of a requirement session persistence was.

[Christian Hopps] (through chat) Each application may have a different definition of "resources" these then have to be boiled down into a single topology Network Aware Computing (NAC! :) does scale

[John Scudder] (through chat) It seems impossible to satisfy that requirement simultaneously with the latency requirement

[Joel Halpern] (through chat) @chris unless I have confused myself, viewing this as an overlay service that uses encapsulations (LISP, SFC, NVO3, ...) allows one to decompose and scale suitably. The ingress edge can include load balancer logic.

[Lars Eggert] (through chat) @dirk exactly. but lb vendors have no incentive to interop?

[Christian Hopps] (through chat) damn lost the AV again.. brb

[Zhenbin Li] (through chat) I do not know how the application layer or the transport layer to learn the network status to steering traffic. For example, the packet is forwarded in the network take the shorted path. There can be other path with light load. How can the application/transport layer detect it and make the network steer to the path?

[David Oran] (through chat) Another fairly deep question is whether the interests of the organization deploying the application and the organization providing the network connectivity are aligned. Google doesn't worry about this because they are both.

[Tony Li] (through chat) I agree with Joel

That's why all of the UPFs have to have common state.

[Cullen Jennings] (through chat) where did Jeff tell us to look ? I missed it

[John Scudder] (through chat) What did Jeff just cite as a place where the problems have been solved?

[Dirk Trossen] (through chat) @Tony so all LBs are same vendor -> vendor lock in?

[Joel Halpern] (through chat) Jeff is pointing at ALTO.

[John Scudder] (through chat) thanks

[Tony Li] (through chat) No, the point is that we need a standardized LB protocol

[Christian Hopps] (through chat) Ah, the overlay is an interesting idea, that is using networking as the LB transport I guess.

[Dirk Trossen] (through chat) @Tony my point of the second question?!

[Christian Hopps] (through chat) The key thing I have a problem with is the current set of drafts I'm seeing where the compute resources are being shoved into the underlay networking layer.

[Zongpeng Du] (through chat) After this layer3 Load balancer, the packet will be put into a tunnel to the Egress PE. So the solution can scale IMO.

[Tony Li] (through chat) @Chris and that's definitely a mistake

[David Oran] (through chat) @Zhengbin: isn't what you describe exactly what Google's B4 does? (<https://research.google/pubs/pub47191/>)

[Christian Hopps] (through chat) @Tony, right, I think my zeal for rtying to shut that down may have blinded me to some more novel ideas like using an overlay as LB transport.

[Dirk Trossen] (through chat) @Chris "The key thing I have a problem with is the current set of drafts I'm seeing where the compute resources are being shoved into the underlay networking layer." I am not entirely sure what this means. The decision of which compute resource is being utilized is being shoved into the underlay networking - is this what you mean?

[Tony Li] (through chat) @Chris Yup, one overlay plane per application. Resources/metric specific to the plane.

Then LB state is distributed so that you have session persistence across the front end.

[Dirk Trossen]

@Tony application or application category, i.e. could we have see the overlay as not necessarily service-specific but specific to categories of services?

[Wim Henderickx]

if you look today at 3GPP and URSP it solve this based on UPF selection. It uses both endpoint + application. Android has a prototype that supports this and was demonstrated in MWC.

[Tony Li]

Not a problem as long as they are willing to share policies.

Dirk Trossen

@Tony Agree!

[Pete Resnick]

More and more, this doesn't strike me as a routing problem; it's all service discovery that can be done in higher layers.

[Christian Hopps]

It still feels a bit off to me, and that the LB as part of the application itself is superior (part of the distributed application itself is to obtain and keep updating the "best" unicast location to use).

[Joel Halpern]

@Pete if not "all", very close thereto.

[Richard Barnes]

@Pete anycast does exist, and is used

[Wim Henderickx]

<https://cloud.google.com/blog/topic...google-android-enterprise-and-cloud>

[Aijun Wang]

The network protocol is used to find the optimal destination. Why not use the more useful information to make the decision?

[Dirk Kutscher]

IMO, the point is that it can be done at different layers, and this is a proposal to employ the network/routing layer.

[Pete Resnick]

@rlb: It's not the anycast that strikes me weird; you could plausibly do SD over anycast. It's the construal of this as a problem for routers that I find odd.

[Tony Li]

The decision that you're making is not just an optimal path.

[Zhenbin Li]

@David Oran Yes. Similar as the B4. If the operators can learn both the application load and the network load, definitely they can achieve better traffic steering result. But B4 work maybe not standardized.

[Tony Li]

You want application information too, so you can't just make a decision at the network layer.

[Pete Resnick]

@Tony: +1

[Richard Barnes]

The thing i find odd here is the idea that you can only move your services around as fast as you can update the routing plane. which comes back to @Pete's point about service discovery (waiting for convergence/distribution as opposed to just updating the SD server)

[Cullen Jennings] (through chat)

@pete - yes - and in some (many) uses cases, if the applications can get a handful of possible services it could use, it can connect to *all* of them and test which is best this make a decision about what is optimal at an applications level. There is a big differences in requirements between "my first packet needs to get to optimal places" vs "after sever round trip to multiple places, I get to a good place"

[Pete Resnick] (through chat) @rlb: Yes, better put than I did.

[Christian Hopps] (through chat) if you want nice things to work like switching from one site to another based on current resource distribution, then the application needs to either be "connection-less" or has to be aware of the switching, and this is not different from being network aware I think.

[Wim Henderickx] (through chat) the dyncast proposal has to work behind or within UPF and in essence the d-router is a load-balancer as Shraddha presented

[Zhenbin Li] (through chat) @Tony the usecase is that the operator can hold both the network and application load. The network can learn the application load for better traffic steering. Yes, the controller should learn the application load info. This the work of dyncast.

Pete Resnick

@Cullen: Are there use cases they're looking at where the first packet has to get to the right place?

Rich Salz

yeah, VR and traffic seems to need the right answer first packet.

Christian Hopps

So what I keep thinking is that someone smart needs to write a framework for L7/applications that provides this network aware load balancing

Wim Henderickx

so another q is. We talk about initial selection, but what happens when the user moves? Do we want to address this? If so we also need to move application context.

Richard Barnes

@Rich quite the opposite. real-time stuff is an ideal position to adaptation

Tony Li

@Robin Yes, but once you've made that LB decision, it must be distributed across UPFs, so there's a lot more to be done.

[Dirk Trossen] For P9, are you proposing L3 LB as in Dyncast D-ROUTER? and for P11 if you have LB at every UPF, signalling needs to handle the coordination of every UPF layer . How do you see the D-router and the controller LB is done?

[Jeff Tantsura] Please refer to (many) RFCs talking about load balancing. The terminology here (?) is not that accurate.

[Dean Bogdanovic] For Individual enterprise and service provider, the balancing is not done in an equally way.

- Solution with protocol extension: (Cheng Li, 15 minutes) 90/120

Architecture of Dynamic Anycast

[David Oran] (through chat) What's the assumed scale of a D-router? 10^6 sessions? 100^8 ? What's the assumed update rate? !Gb? 1Tb?

Typical L4 LBs handle 10^6 - 10^7 sessions

Luigi Iannone

@Pete: the point is to let the network layer decide where is the "right" place, the decision is made by sharing knowledge about the different places

Tony Li

@Wim, my understanding is the state is replicated so that there is no change on a move.

Dirk Trossen

@Wim if your service instance selection does not change, users can move. But yes, if you want to assign new instance mid-session, there needs to be extra work, not at the routing level I would guess.

Richard Barnes

@Rich multimedia apps typically attempt multiple media servers before settling on the one that looks best, as @Cullen says. VR isn't really any different

Rich Salz

yeah, i know we do happy eyeballs racing all the time.

Richard Barnes

ah that's interesting, i was thinking that CDNs / TCP-based stuff might have more first-packet requirements. so it sounds neither case really does.

Wim Henderickx

@Toni, agreed but this is an implicit assumption and need to be aligned with the app.

Daniel Bernier

so client-apps are pretty good right now to know and adapt best edge server based on app metrics. However, to adjust their UPF anchoring is not trivial

Wim Henderickx

@Dirk, indeed my point is you cannot do this in isolation at the routing level. The application part is very critical

Christian Hopps

@joel +1

Rich Salz

@rlb: no, we do both. we want to get the client to the right edge server immediately and then we race to find the best way (or even ip4/ip6) to get the client to the right server

Mohamed Boucadair

all good and fair comments from Joel.

Dirk Trossen

@Wim I do agree, hard to disagree. I don't think the proposal here is that ALL those aspects are solved at the routing level though

Wim Henderickx

@daniel, if you look at the URSP thing in 3GPP the UPF selection is becoming much more dynamic and in my view addresses this. No matter what we do here if the scope is mobile, the UPF anchor is critical. Btw URSP is not the UPF selection but also influences radio resources and schedulers, etc. So it is even bigger than the edge selection. sorry meant not only UPF selection

Adrian Farrel

Erik just made some really key points. There is a very big difference between micro-computations needed in pseudo-realtime (such as in automotive?) and larger computations

Srihari Sangli

Clarifications on the computing resource, its requirements and characteristics would be helpful

[Dirk Trossen] (through chat)

@Adrian +1 but I also think we need to discuss more the basis for decision, i.e., the metrics. Even the term LB is putting too much emphasis on the idea of regular load updates. Is this the (only) approach we are thinking about?

Open discussion -(all, 15 minutes) 105/120

[Lars Eggert] (through chat) i think lb is really the wrong mental model here. the point of lb is to *balance* a large inbound request load over a tightly-controlled and regularized network to a set of identical servers (and at minimum resource cost.) the uses cases here seem to be much more about bespoke "optimal" placement of individual requests.

[Tony Li] @Lars How is that not LB?

[Dirk Trossen] @Tony load may not be the only metric maybe?

[Lars Eggert] in lb, you don't usually care which server instance (and over which path) you land on. that seems to matter a hole lot here?

[Luigi Iannone] Yes Lars, that is the difference in a nutshell

[Tony Li] Granted, there's more than load in the decision. The point is that you're making a decision on the front end.

[Dirk Trossen] in edge computing, HW capabilities, for instance, may be highly important. It goes along Lars' "set of identical servers' point.

[Tony Li] And it's not just a routing decision.

[Pete Resnick] I think to the "questions": It sounds like there are some important use cases, but it sounds like we need to spend more time separating them out in order to answer the second question: There probably are existing and sufficient solutions for some of them.

[Luigi Iannone] @Tony: "informed" decision in the front end...

[Rich Salz] it would be good to have a different term for this use-case than "load balancer" more like "load director"

[Ketan Talaulikar] @Dirk how about we start simple and with just one metric "load". Can that be specified and fleshed out first? Then we know what is possible and what makes sense to tackle in the routing protocols. We can get to combination of metrics as a further step. Makes sense?

[Dirk Kutscher] It's effectively LB, but it's doing more than ECMP....

[Daniel Bernier] @Wim thanks, confirming my point, a DNS/LB feature happening AFTER the UPF selection would be already *influenced*

[Jeff Tantsura] @tony - aren't you waiting for an imminent proposal to flood it in IGP ;-)

Peng Liu

The use case has two parts to get two but seen as one conclusion, steering traffic considering both network and computing resource status.

Ketan Talaulikar

@jeff ... i think we already got one?

Tony Li

Yes, then I roll out my dump truck. :)

Dirk Trossen

@Ketan fair point and something that sounds reasonable, indeed.

Julien Maisonneuve

On compute metrics collection DMTF has a head start.

Jeff Tantsura

yes, a composite/normalized metric would be very helpful indeed

[Luigi Iannone] (through chat)

I think that the routing protocol is just away to distribute the knowledge that the LB++ (in lack of a more generic name) will use to do informed decision is not something that necessarily will impact each and every router in the network

[Pete Resnick] (through chat)

Loops?

[Tony Li] (through chat) If you start shoving other metrics into the routing protocols, then you are very much impacting each and every router.

Christian Hopps

Luigi so that's the dump truck thing.. i.e., don't turn OSPF or IS-IS into a dump truck for distributing state.

Jeff Tantsura

@pete - tunnels solve this one

@chris - sure, we have got BGP for that ;-)

Joel Halpern

@Luigi a) If the proposal said it would encaps, that might be true. b) burdening all the routers in the domain with propagating information needed by ingress edges seems an odd way to approach the problem.

Luigi Iannone

I do not think that loop problem will be worse (but metrics need some thoughts)

[David Oran

Don't we also have to consider the question of "where do I instantiate the computation based on where the network can deliver the data"? This seems to assume the computations are ready to run and just telling the network how to deliver the data. That's the essence of "joint optimization", which can't be done in L3 alone since L3 can't cause computational resources to appear on demand.

[Luigi Iannone] (through chat) @Joel: Yes to both points

[Pete Resnick] (through chat) I'll chat offline. I'm having a real problem getting myself to see this as routing. The existence of "loops" doesn't make any sense to me.

[Shraddha Hegde] (through chat)

whether to bring compute awareness into routing layer or bring network awareness into LB layer is something that needs to be debated

[Lars Eggert] (through chat) @daveo yes. this problem is significantly harder than lb

[Luigi Iannone] (through chat) @Chris: is not about dumping anything in the routing layer ... I think this is part of the work to be done, what to put there and how to use it.

[Tony Li] (through chat) @Shraddha Only if you think that adding more data to routing is a sane thing to do.

[Joel Halpern] (through chat)

Between existing network exposure work, and existing data center exposure work, it seems that the placement problem is an application space problem currently ripe for exploration and differentiation.

[David Oran] (through chat) @Lars" this is the stuff I was referring to when I brought up Sapphire and Ray.

[Dirk Kutscher] (through chat) IMO, this work needs to be informed about how relevant distributed computing systems work, i.e., systems like Ray

[Luigi Iannone] (through chat) @Tony: Yes, if you put information in the routing layer every router is impacted. The more correct statement would be "not all of the routers use the information" which may lead to too much overhead...

[Jeff Tantsura] (through chat) @Luigi - that is why we have overlays

[cabo] (through chat) <https://jabber.ietf.org/jabber/logs/can/2022-03-22.html>

[Zongpeng Du] (through chat) The computing information can be processed just like the OAM information in the network

[Luigi Iannone] (through chat) @Jeff: right. That is one way to do it.

Wim Henderickx

does IETF has a definition or has there been work done on a load-balancer?

[Zongpeng Du] (through chat) only the PE node need to process them

[Jeff Tantsura] (through chat) @Luigi - only the end points (head-end/tail-end) are concerned with the state related, it has nothing to do with the rest of the network

[Joel Halpern] Clarify the proposal, is it proposed to encapsulate?

[Cheng Li] We are not proposing encapsulating the packet. It could be a tunnel.

[Joel Halpern] we have multiple upper layer application. Do we have additional needs for routing (e.g. WG?) or we are using those applications and won't need such new WG?

[Cheng Li] Noted, and will take care in later version.

[Zongpeng Du] LBs (D-ROUTER/D-MA) are good solutions for anycast. These LBs can obtain network information as L3 LBs

[Daniel Huang] For Dyncast, regarding metric service instance, the time-span varies from minutes to milliseconds. When the time is short, is it possible for the router to decide the routing? It is too fast.

[Cheng Li] This is a valid comment. Millisecond would be quite challenging, we need work in an engineering way.

[Jeff Tantsura] How to effectively compute paths? Shall we put CPUs into account?

[Cheng Li] Prefer to go to the list, a big question.

[Erik Nordmark] Regarding the flow affinity, is it from network perspective or from application/computation perspective? Things can be moved and the dynamicness is not only changes in the computation resources but also the distances/positions. Clarify, the requirements to differentiate from CDNs.

[Zhenbin Li] To clarify the B The operator can hold both network status and the application status and thus achieve better traffic steering. This should be noted as the scope of the CAN work. Currently the solution in CAN is pending, however we cannot decide whether the LB must be a part of it.

[Cheng Li] We actually only provide solution for one use case, i.e., the operator is having the network and the computation resources.

[Linda Dunbar] We should focus on the use cases presented.

[Yuexia Fu] The key is a joint optimization. The dyncast provides a potential solution. Given the flow affinity, the dyncast can select the right destination. We can go further in that direction.

[Ketan Talaulikar] To add more with Jeff on Path computation. It's not clearly specified yet. How the computation is defined and measured in routing. Probably rtwg may be a good home for routing-related problem.

[Cheng Li] Agree.

[Dean Bogdanovic] People are bringing similar problem in different scenarios (refer to satellite on Monday). There are more mobility and our architecture may need to be changed accordingly. We need to go to the deep dive and understand the fundamental problem to solve, from routing perspective.

[Greg Mirsky] Something that IETF may work on is latency-oriented. It is hard to use existing work to measure the computation, but we can optimize the latency through the performance monitoring. We have performance/measurement matrix over there. It may be

[Acee Lindem] Joel/Jeff raised the good point. People may start with working from different details but we should aim for converging into a common solution and architecture. It would be useful if people start with the same metric.

[Yizhou Li] Some hint based on current implementation with preliminary idea on how to construct the metric. The key question is an optimization between the computation matrix and the network matrix. Computation can be measured in CPU cores, and anycast IP address can be used to help.

[Tianji Jiang] Regarding the 5Gs, it starts project to expose the computation metric to external and allows integration. So this work is important.

Hard Questions - (Chairs, AD, 15mins) 120/120

[John Scudder] Good discussion and conversation, and it's a successful BoF. It's nice to see the energy, and there are comments/homework to handle offline. Please follow-up on the mailing list. Regarding the questions on the slides, the use cases are agreed to be important. People did mention to look into existing solutions, e.g. ALTO. Please continue pay attention to reusing the existing work. Regarding the proposed architecture, need to double check if put or not put it in the underlay. Load Balancer/D-router or Dyncast has their own characteristics and could be further investigated. Some people in the Chat window stated that maybe we need work to standardize the LB as they are secret sauces. Dirk mentioned that D-Router has the components of front-end LB. Maybe the work is about metrics and messages among LBs.

The next hop of this work (WG to land in) is still not clear, please continue working.