Monday, 16 July 2018

15:50-17:50 Afternoon Session II; Room Name: Place du Canada

RTG rtgwg Routing Area Working Group

Chairs: Jeff Tantsura, Chris Bowers

WG Status Web Page: http://tools.ietf.org/wg/rtgwg/

1) 15:50-16:10 - WG Status Update
Jeff Tantsura, Chris Bowers
20 minutes

2) 16:10-16:25 - draft-mirsky-rtgwg-oam-identify
Greg Mirsky
15 minutes
Jeff: you will be presenting this draft in 2 or 3 more working groups, how do you see this draft progress?
Greg: I have a concern with current state of identification for OAM. The positive of output of this draft would be that there are certain updates to documents in SFC and vo3 groups.
Jeff: after you present anywhere else, I'd recommend you update the draft.

3) 16:25-16:40 - draft-templin-atn-bgp
Fred Templin
15 minutes
Jeff: Thanks for the authors on working with chairs to bring the document to this state. we'd like to it to be adopted at RTGWG. We'll start the WG adoption call after IETF.
John S: what's the status of this doc?
Fred: informational.
Jeffrey: it's an architectural document, informational.

4) 16:40-16:55 - draft-ding-rtgwg-arp-yang-model
Robert Wilton
15 minutes
Jeff: from our perspective, draft is ready for adoption. we'd like to start yang doctor review and an early review from routing directorate. We should try fast track it to next step.
Acee: this is procedure mater. I see you put a number of drafts in adoption tracker. Are you going to send emails?
Jeff: they're in the process of being reviewed by Yang doctors. They are in the process of being adopted but not adopted yet.

5) 16:55-17:10 - draft-ietf-rtgwg-policy-model
Yingzhen Qu
15 minutes
Jeff Hass: I do owe you feedback and it's on my queue. For "set-metric" and "set-preference" you put uint16 and uint8, is there a reason why you didn't go for allowing 32 and then allowing implementations to put modifications?
Yingzhen: For example, the preference right now most implementations you have defined between 0 and 255, that's why we defined uint 8. If you have different opinions we can talk about it.
Jeff H: our implementation goes wider than that. General suggestion, we should define larger range and allow implementations to do variations in YANG through extensions.
Yingzhen: so you are saying we should change the preference to 16?
Jeff Hass: or maybe 32 and 32, and allow it to be restricted. I'll send details offline.
Acee: I believe most implementations have metric larger than 16. Or maybe to have a deviation into a refine statement that refines the range.
Yingzhen: or we may do a union so you'll pick that one that's right for you.
Acee: that would be preferable to implementations that have the smaller range.
Jeff T: so, policies are fundamental for any routing protocol and we are in routing working group, please review the document and we'd like to progress it. There are number of things in protocol that are pending just because of distribution and other things. Please review and see if it works in your implementations, and we will try to find a common way of treating deviations. Please do review it and it's important.

6) 17:10-17:20 - draft-asechoud-rtgwg-qos-model; draft-asechoud-rtgwg-qos-oper-model
Aseem Choudhary
10 minutes
Jeff: do you feel the draft addressed different implementations? I
Aseem: we have incorporated different proposals.

Jeff: any questions? During today's meeting, we're going to have a number of discussions on telemetry and Ignas is going to be speaking as OPS AD. I see a number of implementations of stream

in telemetry and it would be good if you could agree how we represent the data.
Aseem: it's actually not about telemetry, but we can create data set which will suffice the purpose of telemetry. For example, we have timestamp information…(followed by a few examples)
Jeff: There are a number of native models supposed to do this and there are somewhat different than we do here. So it would be really great to converge on how we represent data and progress on this work that could address our common implementations.

7) 17:20-17:35 - draft-lee-rtgwg-actn-applicability-enhanced-vpn
Young Lee
15 minutes
Acee: all these related models are presented in TEAS ACTN, this is in the context of telemetry discussion, that's why this one is here?
Jeff: there are service model and device models scattered across IETF, so my intention is to provide vertical review from service request to device model, from service request to network.
Acee: you don't have the context, it's really at superficial level that I'm seeing this.
Young: the ACTN Yang models are progressed in TEAS wg, but this is VPN+ model, and we see it fits here as information model not as a standard draft. It is how a system can be applied to VPN+.
Lou: I'm ok with it being socialized here, and it's helpful for more people to understand it. Stewart's document is moving to TEAS even though he didn't republish it yet, and this should go with it. I'm ok with having it talked about here because it hits multiple areas. It's a bit redundant, but not necessarily a bad thing.
Jeff: if you look at the workflow of this meeting, next presentation is going to be talking about streaming telemetry, and Ignas is going to talk about management and intention.
Robert(cisco): are your aware of the works being done in MES regarding the life cycle model that seems to be somewhat overlapping with different service layer?
Young: yes, I'm familiar with SLOs. I identify we have this mapping per se, they have a lot of service models and they actually use some other traffic models from other SDOs, not necessarily have their own.
Robert: just want to check
Greg M (relayed): how does the model control performance measurement protocols?
Young: what does it mean performance measurement protocols?
Greg: OAM.
Young: we're independent of it. There is a slide which I haven't explained. Customers want to see the subscribed performance

metric that they are interested in. for instance I want to have strict latency 50ms end of end, but point one to point three I have a different requirement, that's called virtual network. Network collects performance matric that's meaningful to the customer. Depending on the performance.
Greg(relayed): what produce these metrics?
Young: customers are interested in E2E. you have to stitch them together. This will be presented at TEAS tomorrow with more details.
Jeff T: if there is any question, please take them to the list. Maybe TEAS, not routing. Thank you.

8) 17:35-17:50 - draft-song-ntf-02
Fiaccola Giuseppe
15 minutes
Robin Li presented.

Greg: when you say telemetry, do you mean telemetry data as network state information or protocol to collect and/or transport network state information?
Robin: both.
Greg: would you agree that the machine readable is what you refer to as telemetry?
Robin: yes
Greg: just want to make sure you're familiar with L map WG, particular information and Yang models.
Robin: yes. This is definitely taken account, the yang models.
Benoit: in your draft, you said the network telemetry demands new protocols. What do you need we don't have already? What do you mean? What new protocols?
Robin: this slide shows existing protocols/work being done in IETF.
Benoit: for BMP, what you improve is to push it via telemetry. what are you using? What needs to be done here?
Robin: is this new or existing protocol? We're trying to consolidate the work, may enhance existing protocols. We're not going to create new protocol.
Greg: it's critical to differentiate telemetry information and protocol, and scope work properly.
Robin: we'll refine the terminologies, identify scope of different solutions.
Donald: did you try to socialize the idea with potential users? Especially operator community? Whether your solution fits in?
Robin: yes. SPs and OTTs. We've done presentation at MPLS congress in Paris. There are operator co-authors.

Jeff: it sounds like we should jump higher, I'd like to see a proper comparison between existing tools and what you think is missing before telling everybody we need something new.

9) 17:40-17:50 - Coordination of Network Management Evolution

Ignas Bagdonas (OPS AD)

10 minutes

Ignas: open discussion, just about any comments, any feedback, anything.
Dean: I was part of the DT. My issue is that we develop lots of models without actual implementation. we need people to implement models. Until then we will have experience how htings are going. The other 80000 are still happy with CLI. There's still hard requirement that there has to be a CLI and it has to be what we're used to. Once you start talking to Ops, they still like CLI, only know how to debug with CLI. We need to figure out how to do debugging, operational, not just configuration. We're still building configurational, and we need operational model. and one of the things that the existing vendors have is the translation between the config models and the operational models, to figure out what this configurational model means, and the translation between them that is a major problem. We are still building models in the IETF, here's my config model, here's my operational model, and they have to
figure out how will I transfer it from one to the other.
Ignas: yes, exactly that's a valid problem. That needs to be solved.
Dean: About tooling, there are good proprietary tools outthere, no open source unless the group wants to put some efforts. We'll have to rely on commercial tools.
Ignas: my comment was not specifically splitting open source vs. commercial. It was a general comment, yes, there are tools.
Dean: I disagree with you, but that's fine.
Acee: I'd like Dean to review OSPF/ISIS yang models, and see if they are in the sad state of have ops and config split. I think we've done a good job of putting them together. Aside from what you're saying, do you have some strong people from providers to lead this effort? Not only strong, but the amount of time to devote to it, that's the key thing.
Ignas, short answer is probably yes. That is the reason why this experiment is being tried.
Acee: that will be great, I think it's a good idea.

Ignas: this is not something that happened overnight, this has been socialized for quite long time with the community. Yes, I've been trying to pull in some people then twist their arms in trying to do that. If this goes as a WG or an equivalent, certainly the chairs would not be from IETF community.
Benoit: many points I want to make. First, I think it's a great idea to reach out to operators, I agree with code respect that Dean was mentioning. Just working from Yang models without actually having the tools and the code, the first step is not good enough. We keep telling people need yang. Actually, yang is just a means to an end. What we have to do is to show the operators how they will be able to reduce the OPEX, those 80000 would see what they could save if you got the yang and the right tooling. We reach a state that we got YANG modules that they are not coordinated that's the biggest issue for routing. It's good you mentioned for the entire industry it's not there yet. So it's a good experiment, go for it.
Ignas: Yang is the tool, not the goal.
Jeff T: this is not a discussion coming out of nowhere. We've spent the last 6 months with Ignas and few more people working on it. We would love to use yang models, we understand the value but we didn't have the tooling. It makes some degree our work irrelevant because it can't be consumed unless we solve the issue. Otherwise we should ask ourselves what we're doing here.
Robin Li: this is a good topic for discussion. The first we did huge work on Yang, and it should be appreciated. 2nd based on my observation of Yang being used in industry, to some extend I'm pessimistic. We're entering into a negative circle about it, because the YANG model is not for the application. We have to define private Yang models, after the standard one is available, they may not want to change the existing Yang models. As for open source, it can be applied to limited scenarios. So there is a split again. I also thought about possible solutions, we should set up operation community. OTT and IT companies are facing challenges. Lots of information and tools can be shared. 2nd I want to mention that the NMP will be presented at OPS. we can change existing configuration to control protocols. We can think about some of this work using this incremental method.
Ignas: quite important points. Just something can be done doesn't mean it should be done. We have plenty protocols. We should use solutions existing for a simple reason that there is a much higher probability that operator community is aware of it and how to deal with it than to try to invent something completely new which does mostly the same thing.
Jeff: The most beautiful OSPF model wouldn't make for work in services unless there's a logical that enables the consumer to deploy a service unless all the pieces are put together.

Dean: I have a proposal, if you're proposing a DT let's try to start it and I would even volunteer. Within a certain open source project, we've been supporting Yang models in routing demons, and that's pretty active project. Vendors here have experience communicating with operators what they are looking for and we might run a project, we can come up with some proposals and implement them, and see where it will take us and we will have first-hand experience. We'll be first implementing and will be discussing about the actual model.
Jeff T: great idea.
Ruediger: Dean, are you heretically going back to running code thing?
Dean: yes.
Ruediger: actually, the running code thing that not only has been lost somewhere on the way, also the running code that actually can be used.
Dean: there is a group that meeting at IETF quite regularly. Instead of talking, let's come up with the project we can do in parallel and see what works. There's a reason why I drew myself from several drafts because they started to do and work on things that weren't implemented, and I'm not happy with that.
Ignas: so your offer is noted. The timeframe for that is probably around 103, there are still many things to be clarified. There are other events on operation side need to be discussed. That will likely to happen.
Dean: why 103? Why can't it be done during interim?
Ignas: it's not strictly 103. What I was saying for the 103 time frame, something might happen.
Jeff: let's decouple the events. I think what Dean described we should start it yesterday, right?
Martin: I'd like to go judging trust beyond the few people that have commented on the mic and expressed interests. are there people interested? That could be an option.
Dean: there were about 20 ppl in the room. There is something we're doing on the side of IETF, we're just using IETF to talk about it.
Jeff: we'll start the work much immediately.
Robin: I'd like suggesting: we had pieces of yang work. Maybe we should have a DT focus on use cases to exam existing pieces. So it's easy to talk about operations, because they care about use case and solution.
Ignas: it seems that there is a simple way to get the sense of the operators than trying to do a DT. Sounds like we should talk to operators and ask their problems and then based on those answers to develop the solutions.
Jeff T: we'll start talk immediately after IETF to form the group. We've got a number of volunteers. I know all of us have

our primary job if some of you could spend 5% of your time contributing. I see Donald smiling. It will be useful for the community. We got 2nd meeting on Thursday, looking forward to seeing you again.

Wrap Up

Thursday, 19 July 2018, 0930-1200

09:30-12:00 Morning Session I

RTG rtgwg Routing Area Working Group ; Room Name: Laurier

Chairs: Jeff Tantsura, Chris Bowers

WG Status Web Page: http://tools.ietf.org/wg/rtgwg/

---CUPS slot:
1) 09:30-09:45 - draft-cuspdt-rtgwg-cu-separation-bng-architecture; draft-cuspdt-rtgwg-cu-separation-bng-protocol
Michael
15 minutes
Jeff: any hackson result you want to share?
Michael:
Jeff: what were you doing in Hackson? Writing paper or code?
Michel: code based on this slide.
Jeff: we'll do the 2nd presentations, then questions afterwards.

2) 09:45-10:05 - draft-wadhwa-rtgwg-bng-cups
Sanjay Wadhwa
20 minutes

3) 10:05-10:20 - CUPS architecture discussion/authors of 1) and 2)
15 minutes

Jamal from BNG: what are you looking for? First you said you're not inventing anything, then you said 3GPP already defined everything, are you looking for endorsement?
S: 3Gpp has designed base protocol, for BNG, you need some extensions. You need either something new or change something. We're proposing it to be done here.

Jamal: it seems you copy paste the requirement, 99% percent repeated here, TLV, modularity etc..
Dave Allan (Ericsson): also a long time attendee at the BBF. I'd say the problem space that frames the solution space is bigger than just BNG. We've been looking at fixed mobile convergence, one of the requirements is the existing legacy kit talks to BNG can be integrated with the 5G core. There will be a staggering amount of opportunity either for engineering reuse or engineering duplication across the solutions. I'd suggest to take a look at FMC work and think in this context.
Wolfgang from DT: The challenge here is not only to get the tunnels to the customer right but also tunnels towards other providers. Customers can use our line but be terminated on a different router. You need like terminating identifying the customer. The 2nd challenge is keeping the state between CP and UP accurate.
Sanjay: we are leveraging 3GPP solutions.
Georgios: I'm involved in BBF. There are two types of work being done at BBF. One activity completed is for cloud central office: specifics how BNG can be separated, so there is an architecture you could use. If it's an interface between CP and UP, and mobile convergence are involved, this is ongoing work.
Sanjay: points taken. We don't want to have different solution sets. We try to work on frame work that doesn't mandate convergence but at least allow convergence.
Jeff: Michael, if you need to add something, please.
Michael: I just focus on fixed network. The first step is to define interface to CP/UP separated architecture.
Jeff: there are two places of architecture defined, 3GPP and BBF, why 3rd?
Michael: we don't have full analysis of 3GPP work. But we're trying to see whether we can merge the two architectures.
Sanjay: if you look at DHCP, it can satisfy all the requirements. If you look at hybrid access today, customers have deployed it. Now if you have two flavors of cups, nobody use one protocol to control a fixed interface another for the mobile interface. So that's a point where we want a protocol selection to allow convergence.

Michael: There are two gaps. First we need TCP to do this, and 3GPP is based on UDP. We need more discussions. 2nd attributes requirements are different.
Sanjay: The protocols today is over UDP, but it builds reliability, no head-of-line blocking issue. Details, We can work on later.
George: if we work on scenario for fixed access we have solution. if we want fixed mobile convergence then should do

work in BBF. Something like CUPS can be discussed, but it needs lots of time. Maybe start with something possible, then extend the work.
Andrew: this is an opportunity to do something right, but it's dangerous to reinvent something already solved. CUPS is not new in some environment. If we maintain multiple solutions, it's very tricky.
xx from Huawei: should leverage what's done in BBF, collaborate. BBF has produced popular protocols, especially architecture and requirements. That's the right place to work on. I noticed the first presented referenced central office, it's already published. The 2nd presenter, the reference didn't mention CUPS. My recommendation is to collaborate with BBF, there's ongoing work also, e.g. 5G transport.
Sanjay: BBF doesn't define protocols, it defines framework and architectures. Could CU, they didn't take convergence into account.
Xx from Huawei: that might be a good point. I'm trying to say architecture is defined in BBF, and anything related to BNG should be done in one place. The protocol extension can be done here.
Jeff: BBF is aware of the work here.
Evangelos: I think there is a terminology skew. it looks what you're doing is management plane to forwarding plane, instead of control plane. I understand this is also being done in 3GPP. RFC 7426 defines terminology. Maybe map the terminologies.
Wim: we have a unique opportunity to solve the fixed mobile convergence issue. Rushing into only fixed BNG might not be the right way. I see a gap at BBF now. Maybe we should give a response to the liaison to ask BBF do an architecture, then we do protocol extension here. My suggestion.
yyy from Nokia: all that work is already done, no need to repeat.
Jeff: we're out of time. Talking to BBF is another step we're going to take. People has been working on BBF can help. Authors of these two drafts may work together. Thanks.

3) 10:05-10:20 - CUPS architecture discussion/authors of 1) and 2)
15 minutes

---end CUPS slot

4) 10:20-10:30 - draft-yan-rtgwg-srv6-constrain-analysis Yansen
10 minutes
Jeff: maybe you should present in Spring next time.

Martin: please present it in spring, and maybe 6man as well.

5) 10:30-10:45 - draft-hu-lsr-network-automatic-optimization
ZhiBo Hu
15 minutes
Jacob: what's happening here? You're going to switch traffic, then the path is just going to flap. This kind of proposal comes in once in a while, the trouble is before you bring this proposal you may want to look at what happened before.
Jeff T: Jacob is saying by switching traffic you might congest other link, and you may have to move again.
Jacob: this is a control theory problem.
Zhibo: can we discuss it on the list?
Jeff: you mentioned that you need to look at flow information, not only IP. That means you have to keep state of flows, rather than destination based routing.
Zhibo: if the link is congested, we'll adjust the flow away from the path.
Jeff T: how do you know which flow to move?
Zhibo: if there are two flows, one is 80%, the other is 20%, we can decide to change the 20% away from the link.
Jeff: it's more complicated than presented.
Dean: if you go a few slides back. For example this one, you want to move flow a to B, how do you know what's the link capacity before moving so you don't overscribe it?
Zhibo: we count the traffic statistics.
Jeff: let's discuss it on the list. There are significant limitations.

6) 10:45-10:55 - draft-allan-pim-sr-mpls-multicast-framework
Dave Allan
10 minutes
Tony P: you're suggesting this only applies to SR, right?
Dave: this question came up in PIM as well. you can also image this works in LDP network. There is a possibility to do mpls version as well. SRv6 is a bit more question because it's doing pushes and pops of transit nodes.
Tony P: who has really the charter to hold all these technologies in one pile?
Loa: you're kind of relying on MPLS. Have you considered to send it to MPLS wg and get feedback?
Dave: be happy to, I should have planned in advance. Thanks for the advice.
Jiedong from Huawei: it's not very good to consider in data plane. Should divide into two parts: algorithm, and others.

Dave: so you believe this document should be split into a number of items.
Tony P: I'm confused, other elements of PIM signaling?
Dave: it doesn't use PIM.
Jeff: it deals with multicast. It's presented in RTGWG for your information, will progress in PIM.

---SD-WAN slot:
Jeff: we're about to start the SD-WAN section.

7) 10:55-11:15 - draft-dm-net2cloud-gap-analysis
Linda Dunbar
20 minutes
11:05
Lou: did you look at the complexity of implicit in the tunneling safi? The need for doing a recursive look up insider BGP in order to support it? If you do 5512 encap SAFI versus using the encap attribute which is what's in the draft which deprecates the encap SAFI. If you use encap safi, every route that uses a tunnel when you receive the route you will have to resolve through regular update. You will have to locate the tunnel and resolve it. Are you proposing we go back to that?
Linda: we need specific safi, not necessarily going back to 5512. We're looking at creating tunnels among CPEs ahead of time, so controller can push down policies before the traffic.
Lou: that's 5512 plus 5566? To my knowledge, we were the only ones ever implemented it. It was hairy stuff to do. You have lots of overhead in order to resolve routes. There are timing issues. It certainly works, we have working code, but it's certainly complex. I prefer what's in the draft, it's a lot simpler. In terms what you're talking, it's completely viable to make it work, if you want the code it's an open source. But I'm not sure it's worth it.
Linda: glad to hear it. The purpose of this discussion is to bring up issues from other vendors. If 5512 is too complicated, that's fine.
Lou: it's not the issue of 5512 being complicated. This is one of the cases where when you separate out the theoretical definition from the actual implementation you miss complexity. Both 5512 and 5566 are straightforward. It is the corner case. It's the code that gets really complicated. Whether you use 5512 or a new SAFI, not going to change that.
Linda: so you're saying the tunnel encamp is much simpler.
Lou: the tunnel encap attribute that carries the actual routes is a lot simpler to implement and maintain. The tradeoff is you have lot more updates. One is more verbose from the wire the

other is more compact and harder and more complex in code. We have the complexity and the flexibility in 5512, but based on implementation, we said it's not worth it. If you want to go back, we have the technical specification.
Linda: Thank you very much.
Jeff Hass: 5512 was not implemented because it was decoupled. Your route resolution code need know not only IGP but also tunnel mappings. The Encap draft, we're seeing this used in our own internal implementation is for SR-TE policy draft. It's massively overload the router. It does work, it's painful. I'd rather suggest looking at other technologies rather than going down this path for what you're doing.
Linda: are you saying the SR TE safi is also troublesome, right?
Jeff Hass: the SR-TE policy draft where you use BGP to signal is rather complicated. People from larger data centers are driving this, and the draft history itself show that even they're not necessarily agree on how this should work.
Jeff T: you may want to have this discussion in IDR, and you may get better suggestions.
Linda: actually I'm confused. Maybe we should get together and talk about it.
Jeff: BGP work needs to be done in IDR.
LINDA: understand. It has lots of components, and we want to have a discussion here.

8) 11:15-11:40 - ONUG Open SDWAN Exchange, OSE API Interworking progress
Steve Wood (on behalf of ONUG)
25 minutes

9) 11:40-12:00 - discussion

Jeff: thank you, Steve. It will be useful if we can start define service. Taking intent and passing it down. How service model interacts with network. Considering the expertise here, this is where we might be able to help.
Adrian: RFC 8309.
Steve: who wants to do work? If we were to bring drafts in this area, is IETF interested in seeing? In this WG specific? Will you guys provide feedback?
JEFF Hass: IETF has intentionally avoided API work, we focus more on protocols. YANG has pushed us beyond that. But this is not traditionally the best work for IETF.
Steve: understand. Maybe this is the evolution time that this kind of work can be done.

Jeff: we got a few vendors here, not a lot. I'd like to strat the work, I'd like people express their ideas of what and how.
Wim: as a vendor. I like the idea. Given the expertise that this wg has, and there is a gap between routing and applications. It would be good somehow from orchestration point of view we have a single model to support interoperability. Something like BGP might be a good starting point.
Steve: yes, I agree. There are already number of technologies. SD-wan is an overlay technology. MEF is also doing the work. We want to keep the overlay and underlay decoupled. There's a market that wants to keep it as a pure overlay.
Jeff T: the issue of decoupling what consumers want and what we're getting here.
Lou: it will be interesting to hear feedbacks on some of the service models done, or even augment the models.
Ignas: just a comment on the potential value of this work. IETF is good at working at components, not systems. We are good at technology models, but not the service models.
Steve: I see your point. the question is whether IETF want to change or stay in current mode?
Ignas: there is enough work to be done. The question is whether there's expertise on it.
Jeff: I believe the expertise is here. Something has to be done here.
Ignas: that's exactly the problem. The paper looks beautiful, but is it something that can be used?
Jeff: this is what we're trying to address.
Steve: in the user space, it has changed from the consumption of technology to use cases. MEF is doing it. we don't have a body of doing it. We need to create a body that does it.
Jeff: some SR related drafts are describing the relationship between overlay and underlay. There are works already happening.
Wim: is there something we can help?
Steve: the ONUP has the ability to bring use cases, trying to use to drive the vendors. We can take that and translate into actually interoperable specifications.
Wim: I think it's perfect. Go for it.
Linda: I also go to ONUP. I also think lots of work can be done by IETF although IETF doesn't do APIs. Some of the ipsec, tunnel encapsulation, configurations. I think IETF can do lots of work to fill in the gaps what ONUP wants to do. Maybe a focus group. Should be good for the industry.
Lou: I think it's a good idea, and we would benefit from it. Had we had better understanding, we may decide different decisions in deprecating the mechanisms that Linda wants to reintroduce. Or lower layer, like service level we're missing things. We really need the type of information.

Jeff: next step I’m going to poll the WG and see if we can form a DT. If we find there are enough people who are willing to contribute, we could start working. It would be valuable for SD-wan vendors, larger community.
Steve: if you’re interested in the work, you’re welcome to join OSE. There is a link on the web, please contact me.
Jeff: the initial data model is in git hub.
Steve: thanks.

Jeff: any more questions? SEE YOU IN BANGKOK. THANK YOU.

---end SD-WAN slot

Wrap Up:

Chairs: Jeff Tantsura (jefftant.ietf@gmail.com)
Chris Bowers (cbowers@juniper.net)
Secretary: Yingzehn Qu (yingzhen.ietf@gmail.com)