

Wednesday, March 29, 2017

9:00-11:30 Wednesday Morning Session; Vevey 1/2

RTG rtgwg Routing Area Working Group

Chairs: Jeff Tantsura, Chris Bowers

1) 09:00-09:10 - WG Status Update Jeff Tantsura, Chris Bowers 10 minutes

Jeff Tantsura, Chris Bowers

10 minutes

CB: Note well applies. What you said here and contribute on the mailing list have certain implications, please understand them.

CB: there are general procedures in the WG that we use to poll the authors for the IPR, we do that at the WG adoption as well as at the WGLC.

Yingzhen Qu: My first IETF participation was two years ago in this group, and I'm honored to get this opportunity to serve the community. Thanks.

JT: Agenda review. Charter reminder. New RFC 8102. Document status. Encouraging to do shepherding work, there is a number of drafts that need shepherd write-up.

2) 09:10-09:30 - draft-wbl-rtgwg-baseline-switch-model/ draft-wbl-rtgwg-yang-ci-profile-bkgd
David Black 20 minutes

device profile is needed.

yang is good descriptive language. translated into redfish. profile: a collection of models. family of profiles from ietf and industry. from management point of view, you want something common. Roll data centers by adding racks.

redfish: why a new interface? single common interface. DMTF: many organizations are contributing. server management. Swordfish: storage and resource management.

map the yang into CSDL. transfer yang list into redfish. embed router/switch at the root of the schema. redfish json.

yang ->cSDL->json

David presenting.

David: not wearing a TSV hat, this is to introduce a new work.

Joe White presenting, Redfish for network element management.

Rob Shakir: How lossy is your mapping to the JSON schema?

Joe: It is basically lossless. For CSDL it is completely lossless. We have no way to interpret YANG extensions and we turn them to CSDL annotations. JSON schema is derived from the CSDL.

John Leung: CSDL is far more rich. We decided to map YANG to a rich CSDL language.

Rob: Do you have a spec of how that schema for that mapping looks?

John: Yes. The reason why we have two schemas is JSON schema is not standard.

Jeff T: Would be useful if you sent the pointers to the resources to mailing list.

A: there are examples online. (link?)

snip example: interface.yang to csdl then to json. mechanic mapping. one to one relationship. aiming complete yang structure reuse. container become objects, leaves become properties.

[discussion]

CB: what is the comparison of restconf and netconf? The end result seems to be similar, would be good to discuss. odata seems to be richer. You see that YANG model is a good way to start.

John: restconf and netwconf are perfectly good. This approach is meant to be a part of overall converged infrastructure management. Storage and compute people have way more systems that we do. If the infrastructure is managed one way, how do we add network management to that? RC and NC are good ways but different, and that increases the number of management ways in the industry.

Leung: There are multiple redfish implementations in the industry. Compute is almost done, storage is on its way. It seemed for us to leverage the expertise of the ietf for the networking part. GTT have a greenfish working group which covers more than compute and storage only.

Alia: I am happy to see this work. A couple of questions: As you are looking at the yang models, are there features or challenges that you see when integrating into your work?

JohnL: Some of them are not done, having RFCs would help. Yang in itself has a few challenges in itself when you map it but that is solvable. Mostly the set of models is incomplete and there is a lot of drafts, 100 drafts and a third of them are real RFCs if you look now. We should redo that survey again.

Alia: I am extremely well aware that we have a number of models that are still in progress. One of the challenges to encourage which models to progress rapidly is selecting a set that is useful. But that set may lack one or two that are required. Knowing what you need from the models would be good. Many of those models represent a good common abstraction from multiple implementations and use cases. It may be that it is the case that you have in this context.

John: That would be a really good discussion to have. I appreciate your point, there is no criticism in what we said here. Writing such models is challenging. We started writing a full switch schema from scratch, abandoned later but learned a lot from that.

Alia: Thank you for bringing this work.

JT: Yang is yang, when you must and augment when you can. Having running code is very valuable. Please share your findings on the mailing list.

CB: We use rtwgw as a way of introducing new work into RTG area. Is the general idea of those document to be in the rtwgw?

David Black: Yes. Most underlying yang models that we need are already here.

CB: Just making it explicit that we are targeting it to be here.

Kun Xiang from Huawei: This work has some resources, like restful API??? I would like to point to ONOS project that can generate restful API from model. I want to know the difference between redfish and ONOS approach.

CB: Let's take this offline. They produce REST API but they use a richer schema that covers servers too.

John: There are several ways to do the job. It is a way to coordinate servers and storage. If I have core router environment it may not be applicable there. There is a split in a way how you manage systems, there are many ways to orchestrate and control, and this is one of the ways that matches with compute and storage.

CB cutting the line of discussions.

3) 09:30-09:40 - draft-ietf-rtgwg-enterprise-pa-multihoming Jen Linkova 10 minutes

update of the draft.

multihome: first send to correct isp. 2nd what address hosts want to use.

changes: standard->informational.

questions/comments?

[discussion]

CB: As a coauthor. We wrote this as we got a request from v6ops but we did not have a good framework in RTG to understand the overall end to end context. The problem has been floating around for 15 years. There are other documents in RTG both here and in IGP WGs that work on the source address routing portion of this work. We should not necessarily progress this immediately as we need to get a feedback. This is my opinion as a coauthor and not as a chair. We need to incorporate the issues that may come up in those groups.

Robert R: You said this is applicable to the enterprises that do not have their own address space. Should we be encouraging enterprises to get their own address space?

Jen: Some enterprises would like that, but it would inflate the routing table. I am not certain whether a coffee shop should have a /48. You also assume that they will run BGP, but there are small enterprises that do not run BGP but want to have multihoming.

Robert R: Not each coffee shop has a /48. I am not saying that it is a bad idea.

JenL: You still want to have a local ingress, and you want to use the provider's space.

Joel Jaeggli: The distinction between the entity embedded in the other prefix - it is irrelevant how small the prefix is, whether it is embedded in another entity.

Jeff Haas: If you have an internet wide use case. You have more than one address space internal to your enterprise and you would want the users of the host to use the appropriate address space.

Jen: Controlling the internal address space is simpler.

Jeff Haas: You are bringing NAT next to the host.

JT: We will stop here.

4) 09:40-09:50 - draft-ietf-rtgwg-dynamic-multipath-routing Sanjiv Kapoor 10 minutes

no questions

[discussion]

JT: Any questions. Thank you for sharing the results here, it is a valuable research on throughput growth, would be good to see at the next IETF if you have more data and results.

5) 09:50-10:00 - draft-esale-mpls-ldp-node-fir Kireeti Kompella 10 minutes

Kireeti presenting.

Kireeti: I am not here because I want to be here but because you made me to come here. :-)

Acee: Is this really just a description of a technique commonly known as RSVP-TE over LDP?

Kireeti: No. It's for protecting LDP RSVP-TE. There will be a picture soon.

[discussion]

CB: Is this informational or standard?

Kireeti: I think this should be BCP. There is no new standard here.

CB: So either informational or BGP, whatever. BCP has a higher layer of connotation.

Alia: One of the reasons that this draft is better suited for rtgwg is because of the interaction between the components in the context of LFA and other FRR mechanisms, that's the kind of expertise in RTGWG that the draft will get noticed here.

Kireeti: This is the background why we do RSVP/LDP protection.

JT: It is a new draft to this WG, would be good that operators as consumers of this technology read and comment.

6) 10:00-10:30 - RTG-DT-YANG-Arch YANG drafts updates and discussion RTG-DT-YANG-Arch team 30 minutes

Lou presenting,

JT: Lou, could you also present the either decision from netmod wg? Benoit is in the room.

Lou: if you're referring to topic two, it's not finalized yet.

JT: yes.

Lou: Both LNE and NI models are gated by schema mount. There will be planned virtual interim on logical elements representation. If you are interested in VRF etc. please attend because it's going to impact you.

Lou: is the routing area common yang types model ready for LC?

JT: All the common types from IGPs and IDR are there.

Updating of the common types model.

Lou: a few drafts are referring this document, it soon will become a gating document. is it ready for LC? The authors feel it's ready, we'd like to hear comments from the WG and the chairs.

JT: To my understanding all the relevant comments from other wgs have been addressed. we'd request yang doctor review. It seems that it is ready for wglc.

Transition strategy to the data store.

CB: It would be useful to spend 3 to 5 minutes on the history of revised data stores, is it reasonable?

Lou: We have spent lot of time on it. As that said, you're the chair.

CB: Proceed as you see it fit on this.

JT: There are no requirements to change the existing models, it applies only to the new models.

Xufeng presenting on the revised data stores.

Rob: If we are going to do this, we will end up replicating the containers. Why would you want to do that. It's duplicate effort. This is a nice solution, I like it, but I do not really understand how this is a migration and not an alternative solution?

Xufeng: when revised data store is available, this (state container) won't be necessary. it doesn't hurt anything, but..

Rob: Why do I need to support the revised data stores at all? If you look at the solution that we proposed years ago, it has the containers with the state. This does not seem to be a transition strategy, that's the point I want to make. It's an alternate solution.

Chris Hopps: The next slide may answer that.

Xufeng presenting.

Lou: I think Rob's question was on the value of the revised data store. It is a fair point. From my view we want to be able to represent information from the configuration standpoint and retrieve operational information what happened operationally in a clean way by minimizing the amount of redundant schema. That's what revised data store provide us to do.

Lou: The intent is to derive the benefit in how we design and build the schema. Yes you can question whether revised data store approach is the right approach which is a consensus coming out of netmod. Supposing that we are moving to the revised datastores, what do we do in the interim? Wait for everything to be done? Freeze everything, I hope not. What is the best way to get there? We don't want to stop working while revised datastore is being defined.

JT: This is a dialogue, all stakeholders please be open to communication. Shall we support other options? We're listening.

Alia: The milestone for revised datastores is very soon, it takes large amount of work for the changes, and the requirements are changing. Can we set a target when revised data stores are coming. Knowing that changes are coming has problems, too.

Rob: You are blocking implementing vrfs in schema mount. It does not work in yang1. I am confused by the ietf that you are inventing more features. Two years ago we presented our models how it could work. Revised datastores are not coming tomorrow, it's definitely not coming tomorrow, and the amount of code that needs to be written to support this is nontrivial.

Lou: I think Rob you're arguing you want one of the transition approach, the one with config together with operation state. We have a couple of options top include the state, option 1, 2 and 3 can do that. All of those approaches do what you say, just a different way. Number 4 is interesting - there are a lot of models that do not care what the intended vs operational state is, and the ambiguity there is ok. Between now and the time we have revised datastore, you don't have the operation state, but it doesn't matter because you don't care about those cases. In other cases. The question for wg is what of those options we need to adopt as the convention of RTG

area. we're not saying don't do anything, we're asking which convention we're going to use and implement. Which convention do you like?

Rob: with config and op state?

Chris: so Rob, when you objected earlier, your argument is don't switch to a new version because it just looks like an old one. So let's pick one that look like yours so you don't have to change it.

Rob: how many hours have we spent on talking on this?

Yong Lee from Huawei: Xufeng, I have one clarification question. When we augment TE data model using the revised data store, what's the impact? we have many models augmenting other models, what is the impact?

Xufeng: that's why we're here. Last week we were told to do split tree, so the impact is big. If we do the way as shown here, we need to rewrite the models, that will be a big impact. That's why we added other options. Since yesterday this approach is dropped. so don't worry about that, any other options we don't have to do this, we're fine now.

Lou: Xufeng, if you were to pick, which way is the right to go?

Xufeng: Number 3.

Jeff Haas: The reason that we are having this debate over and over again is that there is no single solution. In some certain circumstance, some solution works better. If you have operational state that is segregated anyway you do not need a separate state, and where the configuration is aligned you may need it. There are several ways to structure things.

Jeff T: my proposal is to finish the presentations, and we can do discussion later.

Acee: There are a number of people writing tools for YANG and there are different opinions on which of those before the revised datastores are better representing their views. There may be diametrically opposed requirements and views.

Lou: as a DT we want to come with one proposed solution for the interim with Routing area.

Chris Hopps: I have seen only one person coming up with a different opinion. I'd like to see different opinions. We have operators driving this, BT and Openconfig. If there are people that have different opinions they need to start expressing their opinions.

Benoit: What is the conclusion here?

JT: There are options and we need to get feedback that are best and supported by the majority.

Benoit: later or Today?

JT: Today they are presented, decision will come later.

Benoit: when will it be done?

JT: we'll work it in our WG.

Alex Clemm: We are facing the same issue with topology models. Based on the discussion we had yesterday, we'd like to go to the revised datastore.

JT: The end goal of this discussion to come with a single solution. Hopefully Rob also could be happy with it.

CB: as more point of view, we'd make recommendation of a single solution, not mandatory. If someone feels differently then more power to them,

Alia: I agree on having the strict rules that ignores the details of the technology is not productive, on the other hand if you let every programmer use their own preferred language, you don't have things interoperate. I prefer to hear - letting engineers pick what they like does not work, If we are able to come to a recommendation that if there are models that do not want to follow the common approach we'd like to see that rational behind.

CB: 90% of people do not care what the direction is, just want one direction. For the other 10%, we can deal with that.

Lou: The thing to avoid - models depending on each other, if they end up making different choices they will not be able to align. At the last ietf we tried to do it across models across the while ietf, we couldn't do it, so let's do it in Routing area. If we do fail to come up with a single solution, maybe for start for topology models we say that we do this, for routing and signaling we do other thing. At lease we have some consistency with our type model.

JT: Lines cut.

Xufeng resenting LNE example.

7) 10:30-10:35 - gRPC - draft-talwar-rtgwg-grpc-use-cases

8) 10:35-10:50 - gNMI - draft-openconfig-rtgwg-gnmi-spec

Anees Shaikh

JT: The following presentations are for information only. Thanks for sharing.

CB: We use rtgwg for new work and for informing on the work appearing that is related to rtgwg but not directly there. the draft we are talking here we are not considering for adopting here.

Anees presenting:

[discussion]

Eric Voit: I like what is going on with grpc, also gnmi. It has a lot of functions that are simpler to do over netwcong or restconf. How do we get the work on subscriptions in the netconf to work with grpc?

Alia: We know that there are different ecosystems that interested in using YANG for interoperability. It is not a trivial possibility to document in a standard stream and reference it. We are talking about it.

Benoit: I agree with you that in the end it is about the modelling and the protocol - you may have different ones. If this for information only, or you would have any interest in the ietf? or is it only in github? We are changing the charter in the netconf.

Anees: I see netconf changing the charter to accommodate other network management protocols. This might fit the bill, however we haven't define gnmi as only carrying Yang, at this time we do not have an intention to standardize gnmi in the ietf. There is openness, but from the openconfig perspective and our personal perspective we need much more operational experience with it. It's pretty new, we need lots of experience. We have a bit of time before we consider for standardization.

JT: Cutting the lines.

EV: I would like to see that we can get a group in the ietf to work on this. Can we get a list of experts on grpc and netmod come together and discuss?

JT: There is such interest and this is the intention.

9) 10:50-11:30 - Open Config - what's done/ experiences; what's working, what's not so; how do we work together - open mic Rob Shakir 40 minutes

Alia: We have routing yang coordination mailing list that it falling into disuse, which would be a good use of the list to discuss this and also implementation aspects.

Rob: We have published BGP model that is widely implemented now. Also the policy model is progressing. We would like to progress those in the IETF if IETF would be interested in progressing them. we're also observing some of the discussions happened, like the config and op state issue. extra library dependency we'd like to avoid.

Benoit: Thank you for this. You mentioned few key aspects that models need to work together. Alia is right that if it is not consistent then it will not work. If the metric is the number of models published it is wrong. The metadata of the models is becoming more and more relevant, the catalog work is important on revising the models. It is not feasible to get the model right from the start, we will need to have revisions.

JT: Thank you for the comments. We are taking them seriously and intend to work together.

Thursday, March 30, 2017

15:20-17:20 Thursday Afternoon Session II; Zurich E/F

RTG rtgwg Routing Area Working Group

Chairs: Jeff Tantsura, Chris Bowers

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

JT: Meeting about to start.

Note well applies.

1) 15:20-15:30 - update on drafts: draft-ietf-rtgwg-backoff-algo/
draft-ietf-rtgwg-spf-uloop-pb-statement

Bruno Decraene

10 minutes

Bruno presenting.

[preentation]

Asking for a LC on these two documents.

[discussion]

Les: Curious if you had an opportunity to do test on the forwarding planes that are significantly different?

Bruno: There is no default value proposed in the draft, it is network sand implementation specific.

Les: It seems that the solution that you delivered here has the most of benefit when you have multiple events i a short amount of time.

Bruno: Not necessary. Starting from the first event we need the same delay.

Les: most of the analysis that I am aware of - the largest contributor is the control plane.

Bruno: We are not trying to make control plane faster. If you want every router to wait for the same time you need a single standardized algorithm.

Les: Topology changes in a relatively short time means that some platforms may still be reacting to the first change when the second one arrives.

Bruno: It is implementation dependent, in some we consider the first one and take into account others.

Uma: How are we taking care of propagation delay?

Bruno: Yes and no. No, because the updates are received at the different time and that initial delay is not available. We use the number of events, back off algorithm can be used too. It is implementation specific. We have tried to accommodate for propagation delay, the initial delay to receive the first change.

Uma: For RIB delay I do not care too much, incremental changes are there. For multitopology environment this will be important.

CB: Please take to the list.

Uma: It would be good to cover these cases in the document.

JT: We feel that the document is getting ready for the LC.

CB: I am the coauthor so I am explicitly recusing myself from the discussion.

2) 15:30-15:40 - update on drafts: draft-zheng-opsawg-network-ai-usecases/ draft-li-rtgwg-network-ai-arch

Zheng Yi/Li Zhenbin

10 minutes

Dhruv presenting.

[presentation]

[discussion]

Greg Mirsky. Telemetry is being discussed in many different venues. Telemetry would not give you per flow information. How can you map the aggregate telemetry data to per flow data? You may have different classes of service in the network, and in order to find which flow causes the problem - how would you find it?

Dhruv: There are various techniques, measurement end to end, for delay the path telemetry os suitable.

Greg: It seems that we have a case of terminology. Telemetry is the export of the data if you do not have direct consumer of that data. This is performance, measurement, let's use the right terminology.

Dhruv: This is information that is sent out from the device,

CB: one minute.

JT: thank you. It is a good overview of what industry has been doing for a while. There should be a more tangible content for next presentations.

Acee: Standardizing use cases may drive standardizing the telemetry aspects.

3) 15:40-15:50 - draft-li-rtgwg-sdni-seamless-mpls-mbh

Lu Kai/Li Zhenbin

10 minutes

Robin presenting.

[discussion]

JT: Thank you.

JT: a comment related on network slicing - resource allocation will be controlled by either the core side, or radio side. It would be interesting to see how they are allocated, not necessary only why.

4) 15:50-16:00 - draft-templin-atn-bgp

Fred Templin

10 minutes

Fred presenting.

[discussion]

Acee: You are using future tense. Is this in the operation?

Fred: No, will be by 2020.

JeffH: BGP uses TCP, TCP deals with QoS and other transport related things. How does that work?

Fred: Transport is reliable ground based network??

JH: ...

CB: Are they running BMMP? Routing show before?

Fred: Yes.

Alvaro/AD: We have a charter for this WG. One of the things is to evaluate new work in this area, the work that does not necessarily fit into other WGs. This WG is for evaluation on what should happen to that work - bof, another wg, nothing, etc. Routing in the DC started a couple of IETFs ago. There was an interim in January where some of the things to be discussed today were discussed there, with a motivation on why it needs to be done. We do not want to spend a lot of time coming to a ring whether one is better than the other, but whether there is interest in the IETF to develop something and use what potentially can come out of this.

Stewart: This work needs to be a study of a larger scale, which is signaling IGPs. There are 5g networks that are building IGP network of massive scale, there are other areas than DC,

Alvaro: This is an important information. Chairs need to take into account potential other users and synergies. This needs to be resolved in the WG first and then reported to RTG area.

JT: There is a lot of commonality network DC and mobility use cases.

CB: Can you also specify where this work is being done, could you provide any links?

Shawn: There may be a different thought process required for the requirements of the DC specific aspects. Maybe there needs to be a dedicated charter that is operator driven that focuses on the DC aspects, interop.

Stewart: I know that this is a demand from people doing 5g, I do not know of any organization that is doing this.

Alvaro: We likely can tweak existing or future protocol based on requirements. First we need to have the requirements, otherwise it is difficult to work on something tangible. Possible outcomes - yes, we might do a new WG. That WG needs to have a charter. That could be one of the potential outputs. Another output - if we focus on solution X, that could be worked in a specific WG. This is the second time that we are having this conversation, even though that this is the function of this WG, we do not need to exercise that much. It is about having the right solutions for developing such solutions. The question whether we have enough people in RTG and IETF to review all possible solutions.

Stephen: There is an order or magnitude more hosts in the network and that might be a starting point. We need to think about the incoherent operation of ten systems, it works well as long as it is coherent. We designed things 35 years ago and now we have much better understanding on many aspects,

Alvaro: Chairs need to coordinate this and come with an answer.

5) 16:00-16:10 - ISIS in Spine-Leaf - draft-shen-isis-spine-leaf-ext

Naiming Shen

10 minutes

Naiming presenting.

[discussion]

Shawn: Do you have any fragmented or disjoint CLOS topology in your slides?

Naiming: No.

Shawn: If spine loses connection to one of the edges, does it stop sending to everybody?

Naiming: We will cover that in other slides.

Shawn: Sending default in the fragmented topology - you either have all the routes, or miss some, as you cannot have a default any more.

JT: 2 minutes.

CB: We have used the question time.

6) 16:10-16:20 - BGP SPF - draft-keyupate-idr-bgp-spf

Acee Lindem

10 minutes

Acee presenting.

[discussion]

RobertR: In modern DCs compute nodes are L3 devices, and I am running routing to compute nodes. How in this model am I supposed to advertise reachability?

Acee: Do compute nodes have full set of routes?

RR: They advertise a set of routes, or could advertise a default.

Acee: They do not need to be part of SPF underlay that would be a form of service.

RR: They can be part of overlay but do not need, solutions like Calico are attractive just because of that.

Acee: Running dynamic routing protocols just to inject to the flat network - yes you can do that as well. It would be like a route redistribution in igp.

Uma: One of the changes is to flag the SPF run?

Acee: Yes, just to determine that you have the up to date versions.

Uma: This is the fundamental change to BGP, maybe you may want to use a different port? for some address you are using bestpath, for others you do not.

Acee: You deluxed based on AFI./SAFI value.

Uma: Russ's proposal is very modest, maybe you need a different port.

Acee: We separate by address family.

IgorG: Trying to understand =- the difference between traditional CLSO architecture, are you trying to minimize the amount of sessions from each node and converge faster - is that the only advantage compared to traditional BGP full mesh?

Acee: Yes. There are others too. You need to be able to do fast rehashing of your ECMPs in your CLOS topology. BGP SPF is targeted to DC topology, later it can be applicable to other deployments too. Microloop avoidance is still a consideration even in CLOS.

Igor: If you design it correctly it is not a problem, there are two ways to design it, one that has this problem, the other that does not.

JT: Would be good to provide the pointers,.

Igor: It seems that you are trying to solve a problem that does not exist. I cannot tell you the exact topology but may show the device utilization.,

JT:L Would be good that the author group comes with a solution proposal.

7) 16:20-16:40 - Open Fabric - draft-white-openfabric

Russ White

20 minutes

Russ presenting.

[discussion]

IgorG: This works for CLOS topology only.

RussL: It has to be 5 stage. Draft talks about 3 stage and you need to specifically configure it.

Igor: Many of us have done such topologies since 2011. Tor topology.

Russ: I am not certain whether butterfly type of topologies would not work.

Igor: I am not certain whether we need to solve yesterday's problems tomorrow.

Russ; Hypercube topologies - that is interesting, let's talk offline,

RR: I have the same question as to Acee - how do you advertise dynamic reachability to servers in this case?

Russ: I will run dynamic reachability to the servers.

RR: Would be good to have uniform end to end story.

Uma: I like the part where you do the reroute. Minor comment - what you do if you have parallel adjacencies?

Russ: Valid comment.

Igor: There is a chipset modification that was required to use the alternate forwarding modes, but now when that modification is about to ship it is possible to do that. In my CLOS design it is not multi tier, it is multiple dimensions. Use it, then have a clos of closes.

Russ: many years back we did a PoD design using hypercube.

Shawn: in addition to what Igor said, it makes sense for us to have a topology in each layer.

Russ: We do not aggregate, we want full routes.

8) 16:40-17:00 - RIFT - draft-przygienda-rift

Tony P.

20 minutes

TonyP presenting.

[discussion]

CB: 4 minutes left.

CB: Alia has asked me for 5 minutes to conclude and provide the next discussions steps on YANG discussion.

9) 17:00-17:20 - Routing in DC - open mic

20 minutes

CB: 4 minutes on DC discussion.

Igor: This is cool work. You mentioned that noone does the same level connectivity - that is false.

TonyP: I said vast majority does not, it is my data set.

Igor: Some people do, short or long lived. There may be topology change cases that requires it.

Tony: I do support leaf to leaf.

Igor: You mentioned everyone goes up. That is not the case.

Tony: That is a misunderstanding.

Igor: We are talking about close that we started using in 2011, many of us are looking into next things, esoteric topologies that eventually will ship. Dragonfly as an example.

Tony: As long as we know what those topologies are then we can try to make it work.

Igor: By the time this is standardized, coded, and deployed we already may be on the other topology.

CB: Would you be willing contribute to the requirements document? We would appreciate your input. Not necessary debate on the list but to guide out evaluation process.

Igor: That would be my requirements.

JT: That would be very valuable. Based on your requirements people can base how other people may base their design on.

Tony: I worked towards what the pain point are, and 100% of pain points are with fat trees.

JT: Igor, please share the information.

Igor: This seems to be an odd problem. BGP works today, not certain whether that will work tomorrow, but we should not optimize for yesterday.

CB: Please specifically describe what you need.

JT: Please also ring in hardware people.

DY/Yandex: it is more about routing in the DC - would be interested to start a group on the requirements and topologies, hierarchical boxes now becoming inexpensive it looks that a need for topologies that are not clos is coming down, you can do 95% of use cases with 2 or 3 spine layer fabric.

JT: request to you - please document the requirements.

Alia: A quick update on context on YANG, One of the questions we had was how should we handle operational state. We have several options for recommendations. We need to nail down the details for some recommendations. This conversation should go on the rtg-yang-coord as well as rtgwg. We have recommendation on having state and config trees, and that is what we do not think how models should look like going forward. We have separate -state and -config trees, that is not the direction. We're not going to block existing work, but we'd like to see them being enhanced. we're going to send out details later today with details. We talked about the old train, we end up with duplicated data. Two trains are coming - train 1: you can sprinkle containers with config true/false noses through the tree. Train 1 means we can write code today. We need implementation experience and see the interactions. Train 2 has the revised datastore dependency. That is the future. We are looking into definitions defined in 2-3 months, but it will be 9 months before it is stable. Operational state is the longer horizon. Some models do not have dependency on the operational data. If your model does not depend on operational data you can go directly to train 2. None of this is you must do it, some of this is we're learning experience. We will get more discussions on the mailing list. Questions?

CHopps: Transceiver power example - if I set it up somewhere, it is not necessary that it is the actual value is the same, I will read it from the actual value.

Kent: 6871bis will be updated. It will be sent to netmod and ccd here.

Alia: we put all this together quick. We'll look at phrases we want to use to clarify.

Xufeng: for train 1, Deprecation is obvious but destructive, we need to put a new revision of the model. We have two other options mentioned yesterday. State containers would be less disruptive.

Alia: I think we are looking into containers for deprecation- one of the tricks with this is that we need to have implementation experience.

CB: Taking offline. Session is closed.

Alia: If there is concern let's talk.

JT: Xufeng presented 4 considerations yesterday on how to proceed, please take a look. Thanks you and see you in Prague.