Network Virtualisation Overlays WG

Wednesday Afternoon session I - 13:50-15:20 – Boromphimarn 1/2

1. WG Status update
(WG Chairs, 15mins)
Agenda bashing.
Milestones updated.
May need an interim meeting.

2. Security requirements update
(Daniel Migault,10 mins)
draft-mglt-nvo3-geneve-security-requirements-04
Daniel presenting…
Two sets of requirements: SEC-OP and SEC-GEN
Do you think the two sets of requirements are valid or there is one we should remove or if you want the requirements to address another type of question or what is the feeling of the room?
Matthew: Have both operational operator requirements so they can judge whether or not deployment is secure and requirements on a protocol to meet a certain security criteria in the same document, seems to be confusing to issues.
Daniel: removing some text?
Matthew: Should you have an operational security requirements draft separate from requirements on the protocol designed to make it secure?
Daniel: Yes, we can split the requirements into different drafts. I'm fine with that we can have one draft with a threat description and one for the operational requirements and another one for the Geneve security mechanism requirements. It's three document instead of one. So I don't know how it's going to be. The overhead or the commitment…
Ilango: I submitted in a bunch of comments on this draft and also I submitted a comment earlier today as well. There is a fundamental disconnect on how to frame the requirements for the protocol as well as for the operational evaluation of the operational requirements. The draft makes an assumption that you need to meet all of these requirements in order to claim that this is a secure Geneve deployment. But that is not necessarily true. Because the operator will evaluate the risk associated with their specific deployment and see maybe one or two or a few of the requirements that they need to satisfy. It meets the objective that the deployment is secure, whereas this document takes an approach that pretty much all the requirements have to be met in order to claim that it is a secure deployment. I haven't seen any precedents for such an extreme approach for other transport encapsulation protocols. So why we have

to go into this level of details? And In fact it gets into implementation details as well in some of the areas which is burdensome for the operator and the implementers. Not necessarily needed. Take any product column and replace Geneve with any other protocol like LISP or MPLS over UDP or IP over IP, none of those protocols have to go through such kind of requirements. But people are still able to securely deploy those protocols and successfully operate it as well. That's the fundamental disconnect here. So I think we need to address that one first, otherwise we have been going in circles for the last two or three iterations.
Once this resolves, we can always go back and fix some of the individual requirements. That's basically my comment. I think it will be good if you can get more input from people who are in the operational experience or the people who have protocol implementation experience, that'll be really good.
Sam: Speaking as individual. One thing I'd like to know from Ilango, is it with respect to the classification that a specific protocol or in this case the security requirements for Geneve indicates whether it is secure or not, or are you worried about the amount of requirements out there…?
Ilango: The concern is for both cases. If you look at individual requirement, I think some of the examples, if you go through the subsequent slides you will see that, where the document talks about multiple different requirements, and also claims that you need to have a Geneve specific or a security mechanism, and then all of these requirements need to be met in order to claim that Geneve can be deployed or it can be used in any Geneve deployments. And that's the approach that the document takes. But I don't think that is necessary.
Sam: If you remove the criteria that the Geneve could be secure and these are the requirements which could be met and we need not classify by meeting them or not meeting them will make it secure. Does that address your specific concern?
Ilango: This could be one possibility. I think we need to first address how do we frame the problem and that's basically the whole challenge. What is the documents objective? Is it to go and give guidance to the deployers that these are the possible or potential risks in an environment? If you have to address risk number A then you need to fulfill these two or three requirements in order for you to deploy in a secure manner. That's how it needs to be approached. By that way it gives guidance to the operators so that they can evaluate the risk and accordingly pick the right requirements that meets their needs. Then you can go with that approach rather than saying that you need to… If you prescribe like ten requirements and say that all these ten requirements have to be met in order to call

that this is a Geneve deployment. That is unnecessary and other people might start to ignore it. We really wanted to give guidance to operators not to over-prescribe something that operators may not really use or care.
Daniel: The ways I try to write the requirements, especially the operational one related, are we have classification for each of the threats we mentioned, and I try to be very cautious in the text to provide a condition when the requirements applies, which means that if the condition is not met, you considered the requirement is …? Requirement applying is a condition match. For example, I'm saying that by default you should encrypt the inner payload, but I'm also saying that you may disable this capacity if you believe so or if you don't believe the risk is too high. So in that case, my intention was to say, if you don't know you encrypt, but if you carefully evaluate that the risk is sufficiently low, then you can disable this mechanism. As long as you check those, I would say security operational requirement one is met.
Greg: I just wonder how can we get their input from operators themselves and not speak on their behalf. So might be some blind poll and ask them to evaluate? Because otherwise it will be just my opinion, your opinion. We need some objective information, not that what we think that operators will do or not do. And then we can just make it practical and realistic. Something that helps operators. There will be certain things that you need to do, you like it or not, if you want to claim that this deployment is secure. I think that we need to seek information from operators directly.
Sam speaking as an individual: What you say is we provide the necessary guidance and let the operator choose. It is pretty much IETF's problem, not just NVO3. Hard to get the operators provide feedback.
Greg: We have some operators here. They might want to volunteer and suggest how they can give us information in the blind poll.
Sam: Feedback is always welcome but we can't force someone.
Daniel: The requirements came from an analysis from the threats related to the Geneva Protocol and its deployment. So we are not going into operational deployments that are not related to Geneve. So I guess at that point the security analysis might be sufficient to provide a complete set of requirements. If we have a list and you say you pick the requirements you want, it means that if you are able to meet one of the requirements and you will be able to claim that your deployment is secure, it's hard to buy or to sell. So this is why I don't understand how we can have a subset of the requirements and people choose what they want.
Frank: I have a more fundamental question, how are you expecting that somebody reads this document given that it's informational. Is it

just like you want to use this document to self assess yourself to understand whether your Geneve deployment is considered secure and what security means, or do you expect that an implementer would follow this. So I'm a little puzzled on what the document wants to go and achieve.
Matthew: This is part of my point earlier on about mixing up the operational security requirements with protocol requirements on the security.
Daniel: SEC-OP, if you have a deployment and you want to check how secure you are, or if you're a secure or not, it provides you guidance. The security Geneve mechanisms requirements, those are intended for architect or protocol designer. And to elaborate what we called Geneve security mechanisms which means it's a mechanism that can secure any Geneve deployment. The latest set of requirements will be used to define how we can define Geneve security options or a way to secure any Geneve deployment. It means that if you implement this one, you don't have to look at alternate solutions.
Matthew: show of hands for who read this draft? One.
Sridhar: I have similar comments as Ilango's. It starts from the basics how to actually frame the requirements and also in terms of the all or nothing approach, so we do want more people to read the draft and comment on the list whether this is the direction that we want to go on.
Matthew: more feedback on the list, discussion on the list. We've been talking about having a virtual interim meeting just on security to resolve some of these issues. We will send a notice for interim meeting some points in the next month or two.

3. Geneve WG LC update
(Ilango Ganga - remote, 10 mins)
draft-ietf-nvo3-geneve-08
Ilango presenting...Update on Geneve WG last call.
Greg: Appreciate your consideration with my comment on O bit. You decided to remove for reference to OAM, a little bit confusing that it's now indicates the control between NVEs. Is it implying that there is an in-band control protocol between NVEs in Geneve? Geneve header already have a c bit, so how this new interpretation of o bit compares to the existing definition of c bit? I thought that c bit is a command bit. This is a control bit. what's the difference?
Ilango: These are two different bits. We did not change the description of the o bit. The description of the o bit remains the same it was meant for exchanging control messages. The idea is that when this bit is set, it indicates that control message is being exchanged and those control messages can be forwarded to an exception queue so that you can process that in a slow path or an exception

path. That was the intent of this bit. Examples of control messages could be like a virtual link heartbeat kind of a message or could be BFD kind of a message. These control messages are different than the data message and needs to be handled in the control plane to control path. The C bit that you talked about is actually critical options bit that indicates that the Geneve actually carries no critical options. There's also a description on what it means by critical options and how do you handle that. That's been clearly described in the draft.
Greg: I want to point that not all oam are processed in a control plane. So generalizing that oam as a control that is processed in a control plane, that is very risky. And will lead to a lot of ambiguity. I'll encourage to just remove this bit from this specification and leave it for later discussion.
Matthew: There is plenty of precedent for having an associated control channel. Pseudo wire for example, it doesn't define whether you process it in their data path or you process it in control plane.
Greg: Agree. Since we have a protocol field, the protocol can define associated channel and there is no apparent reason to have another mechanism to indicate that this is associated channel message. So the associate the channel sufficiently identified by the protocol type.
Ilango: It need not be a separate control protocol. So basically the next protocol could be the same. For example next protocol could be Ethernet and your control message could also be an Ethernet. So this flags that it's a control packet and then it needs to be handled through the exception path in our slower path and that's how it's being used and so that's the reason why we have this bit.
Greg: We're doing in circles because we already have this discussion what's the purpose of using Ethernet encapsulation for NVE to NVE messages. They can have just their control message format identified by next protocol and then similar identifies the type of the message. So you can decode it easily. I don't see any rationale to have this bit used. Using dual mechanisms of flag and next protocol to identify oam creates ambiguity.
Ilango: We will remove the reference to oam here. We are talking about control message.
Greg: We identify the associated channel easily by the next protocol. You have a shim layer that identifies the type of the message and associated channel as we do in pseudo wire VCCV MPLS LSP. There is no apparent need for the flag. Why make it several ways? It creates a problem with interoperability implementation.
Matthew: If you're running VCCV with IP encapsulation in the pseudo wire, then you have both bits indication this is a packet on the VCCV channel, it creates the exception that you then look in it's an ipv4 packet and that tells you then maybe it's a VCCV ping message. You

have a type after VCCV exception. So it distinguishes it from a raw IP packet carried belonging to the user data on pseudo wire. So there are number of precedents are being an exception bit and then a protocol ID.
Greg: It creates ambiguity.
Sam: You can choose…If you do not want to implement it, yes…
Greg: It causes more. You need to advertise you capability otherwise there will be a problem with interoperability. That's exactly the situation I am concerned with. If people choose different not compatible ways to implement it, then oam would not inter-work.
Matthew: I don't think that makes any difference.
Frank: see the updated text and then re discuss? Because that might help clarify whether and how that bit would be used in the future.
Matthew: You need to probably post the new version of the draft. I'll keep the last call open for a while, and I would request people to review the new version of the documents. And we'll give it a week or two for comment.
Ilango: The proposed text was also in the mailing list so you can also take a look at it. We will publish a draft within a week or two that incorporates all the comments.

4. iOAM Drafts relevant to NVO3
(Frank Brockners, 10 mins)
draft-brockners-ippm-ioam-geneve-01
draft-weis-ippm-ioam-eth-00
Tal Presenting…
Greg: We had a very interesting discussion on security aspects in Geneve. So what are mechanisms for integrity and confidentiality protection on ioam over Geneve?
Tal: We have POT being discussed in the SFC working group, but it's not necessarily strictly for NSH. It is a generic mechanism for IOM.
Greg: POT doesn't have protection of integrity of data transported in ioam.
Tal: It has an optional mechanism of protecting the integrity of the IOM metadata.
Greg: Which is?
Tal: It's part of the draft. To have an HMAC which covers the IOM part, and then you use the Shamir's secret sharing for sharing the secret between the nodes along the path.
Greg: So that's integrity confidentiality.
Tal: We haven't analyzed that yet, and actually it's a good question whether we need confidentiality in this case. Do you think we need confidentiality?
Greg: I think it's discussed in context of Geneve.

Tal: Okay so it's a valid question. Any other feedback about that specific question we'd be happy to hear.
Greg: Another question is I understand that now the proposal is to use next protocol code point to identify ioam payload and place it after the Geneve header ahead of payload protocol data. What happens if the node doesn't understand?
Tal: We presented two different ways here of using IOM over Geneve. The first one as we said an IOM blind node can skip over, the second one an ioam blind node can't skip over. That's one of the main differences between these two options.
Greg: Another difference is that TLV is limited space in Geneve, so the amount of information that you can put in TLV is very limited.
Tal: 128 bytes
Greg: Yes, it's limited especially if you add a HMAC on top of it, 16 bytes, so what left and couple nodes.. so basically you'll fill with HMACs and overhead for TLV. My question is so why not to collect and transport data separately of the data packet? So you can hit with IOM data measurement, but then collect data in a separate follow-up packet. And there is already a proposal for that.
Tal: Thanks. Any other comments or questions.
Frank: If we use Geneve, we'd be as much protected as any other TLV in the Geneve header from an integrity and security protection. So we basically leverage what the base protocol gives us. So there's nothing specific. I think the request then is, it really comes down to we want a code point for Geneve to go and carry iom data. I think that's the real base request to the NVO3 working group. The other option, carry it side by side or carry the IOM data along with the Geneve header, that's almost like in parallel. It is something that people could go and implement, but it is not specific to Geneve. It's just Geneve happens to use an Ethertype type. And people can go and sequence the header in as they could do with other protocols say GRE whatever. So I think from a discussion focused perspective, I'd appreciate thoughts on whether people see that as useful as the earlier common. There is a restriction because the length field is just five bits so 128 is relative restrictive. That's sad. I've not seen deployments with Geneve so far, especially in overlays that have a load of hops. So 2-4 hops might be completely feasible. Any other thoughts?
Barak: looks may be theoretical, considering the lengths as a limiting factor, is it on a table to extend that?
Tal: to extend the size of the tunnel option?
Barak: the option area
Tal: Any other comments?
Michael Smith: Will an ethertype already go forward for Ioam regardless of what Geneve does?

Tal: First of all this is a new draft so we haven't done that yet. But one of the things that was discussed in IPPM this week was that one of the steps along the procedure here is to allocate an ethertype from IEEE. And apparently this is a well known process here in the IETF.
Michael: So there is use case beyond just Geneve?
Tal: Right, this is just one of the use cases.
Michael: So if it becomes an option we would also have to deal with the case it is in the inner packet anyway.
Tal: Right, assuming this draft is adopted.
Ilango: Regarding the length of the options, this was already discussed and I know there wasn't a consensus to make that change and also it were shown there are already two different mechanisms. So if something that can fit you, it's not just the way a message alone, you are to carry other options as well and that's the reason for the whole TLVs just not to carry one option. You also have an alternate mechanism if you wanted to have length larger than 128 bytes. As it was shown here you can always come up with different mechanisms for oam in purposes.
Barak: If I get you right, you said there is a mechanism to extend beyond 128 bytes. Can you please explain?
Ilango: I didn't say extend it beyond 128B. Basically what I was saying is one of the methods that was shown here is to use the TLV and the other mechanism is to use the next protocol type in order to carry it. There are two different mechanisms proposed as part of IOM.
Barak: I think we're kind of forced to trade-off capabilities from one end we will not allow device doesn't support IOM to be able to parse, on the other end when we implement it using TLVs we are currently restricted with 128 bytes which could be sufficient for some uses while some may not. If we can combine the benefits and have larger TLV space maybe it's a good solution.
Ilango: I think for most use cases this has been working well, the 128 bytes. And as I said for ioam I think there are other ways to solve this problem and I don't see a reason why that needs to be by increasing the TLV size.
Barak: Which ways?
Ilango: Isn't that that was shown just now in this presentation? There are two different mechanisms.
Barak: From my perspective it's a wrong way. It doesn't allow devices that do not support IOM to process general header and the overlay. I don't see it as a valid solution.
Greg: I probably use a stronger language than Barak. Introduction of ioam in a heterogeneous system will break the services because packets will get dropped. Not just the node that is not supporting ioam in this manner with a dedicated ethertype will not process ioam,

it would drop the data packet, and that definitely should not happen. We need to understand that this proposal requires the system be homogeneous, the whole domain has support ioam. Their way of using TLV is less intrusive. You don't support it, no harm done. You do deliver data, you just don't get the measurement, not good, but no harm done. In second proposal there is a harm.
Tal: So you mean you're supportive of extending the option to over 128B?
Greg: No, I'm not supporting, just illustrating to Ilango that the second proposal is harmful.
Tal: Okay. As Barak said it's the other alternative, since we don't want to limit only to up to 128 bytes, or presenting two options, and I think it's a bit of a broader question not just in this scope but in general, or do we think there are other use cases that will require more than 128 bytes? This was asked in the past but maybe it needs revisiting.
Barak: A little bit disagree with what you said because I think that it's not the way that we should choose between one of them. We should have a good solution. I think none of these currently is good enough. Maybe we can try to make the TLV based better.
Greg: Or just to choose the different method. Collect it in a separate follower packet with no restriction to their length and no risk to impact the data flow. So basically I'm just pointing out that there is a third alternative. Do not piggyback data on a data packet Do not piggyback back telemetry information on a data packet.
Barak: I just can share that at least some physical operators is that duplication and then sending the data in parallel is not a good option from their perspectives.
Greg: I think that we reached the point to have this discussion on the list.
Matthew: One of the discussions the chairs are having at the moment with the authors is where this work resides and be interesting to see if you could raise your hands if you'll be interested in in reviewing these drafts within NVO3… (A few). Okay, let's give this fair number of people interesting this work here.

5. Overlay Subnet Architecture
(Ting Ao, 10 mins)
draft-ao-nvo3-overlay-subnet-architecture-00
Ting presenting…

6. Geneve applicability for service function chaining
(Sami Boutros, 10 mins)

Sami presenting…

Greg: When you receive Geneve packet with the NSH payload and this SFL list, so you create certain state, is this state only persists until you receive this NSH packet back or it persists for longer?
Sami: It depends on your implementation. You can choose to optimize for the implementation, you could have it persist for a while. You create some flow state and you keep the flow state if you are monitoring for example a UDP session, till it times out… So there are options for the implementation, but I don't think we are going to be defining option because it's up to the implementer to do.
Greg: I think that would be helpful just to explain possible ramifications of aging and having stale entries.
Sami: a good point.
Greg: Another question is what happens if behind this NVE you have sub chain. Basically it's not one SFF, but it is several SFFs that located behind this NVE with their SFs attached, and somewhere there you have a re-classifier…
Sami: I think it's described in the document…
Greg: But you have re-classifier. You reclassify your packet and you change the SPI. So basically what happens… because then you return back and you don't have a cache…
Sami: …new classifier… and then it will have that extra information meaning a new SPI…
Greg: So basically we end up possible worst case scenario. If we have a re-classfication in each in every NVE, we need to provision them from the control plane, right?
Sami: Two things we have to differentiate here. This is not redefining SFs…
Greg: No, I understand, but my understanding of use of this SFL is that we only program the classifier and everything else is populated dynamically through SFL.
Sami: What do you mean?
Greg: You use control plane to program the ingress NVE which co-located with a classifier, but if we have and that's SFC architecture permits to have a classifiers or re-classifiers down the SFP, so in most outrageous scenario we have re-classifiers at each and every NVE, so that means that each of them needs to be programmed for the control plane, so basically this is the worst case scenario. If that's the worst case scenario, then why not to use a control plane to begin with all the time? the program, the state and mapping between SPI…
Sami: we are not saying you have to use that option, two options…
Chairs: running out of time….

Frank: So you have another option class defined and you have within that option class the ability to go into sub TLVs, one of your sub TLVs is HMAC and that sub TLV is already consuming something like 40 bytes right, so is 128 enough of a length for you or do we have another use case for the future that 128 again might be a little bit too restrictive?
Sami: Sure that's a good point. I think you are right. I was entertaining the discussions on 128 byte that just happens
Frank: so it's already two use cases in one meeting, not bad.
Sami: We are going to be presenting the draft as well SFC tomorrow.
Sam: You want to adopt it in NVO3, right?
Sami: Correct.
Matthew: Raise your hands if you'd read this draft? Six.
Chairs: Close this meeting.