

IETF100 NVO3 WG agenda

Network Virtualization Overlays WG

Wednesday Afternoon session I - 13:30-15:00 - Canning

Matthew Bocci: Meeting will start in a couple of minutes.

1. Administrativa.

WG Chairs.

5 mins.

Matthew: Meeting starting. Note Well applies. Blue sheets circulating. Normal meeting this time, no roundtables. Agenda bashing.

Matthew: Milestones - they need to be updated. Any comments on milestones? None.

Matthew: Document status. Encapsulation documents would need transport and security reviews. VMM document has been around for a long time, it is time to LC.

Behcet: I am one of the editors of this draft. It has received good reviews. I think it is ready for LC.

Matthew: Control plane drafts. There are drafts targeting the main areas of control plane. NVE-NVA will use YANG and related transports for communication. At the moment only VXLAN YANG module is available.

Greg Mirsky: I have looked into Geneve and what would be needed for YANG work. The definition of some common TLVs would be helpful for model development. Are there such plans for defining common TLVs?

Sam Aldrin: Have you talked with authors?

Greg: Not yet.

Sam: Please discuss with them.

Greg: I believe that there is benefit in having common TLVs.

Matthew: Applicability of EVPN as a control plane mechanism. The charter is a bit strict on the control plane solution. We cannot talk about BGP in this WG.

Matthew: Security drafts - there is some progress. We need very strong security guidance for encapsulation. The guidance that we are getting from our AD is that we need a normative reference to security drafts. At the moment those are informative. We need to have concrete progress on this topic to get Geneva moving.

Matthew: OAM progress. Over the years there were multiple OAM solutions being discussed in NVO3 for different dataplane encapsulations. We try to adopt more generic drafts too for encapsulations. How should we move forward with OAM - this is a question for the WG. It seems that there is no consensus on a common new generic OAM solution.

Greg: What is important - to differentiate active OAM and other methods, IOAM and AMM. Hybrid methods exist when there is a probe packet. If we consider that we need to do service activation testing, to do verification of service without active traffic, then active OAM is a requirement. If we think that it is insufficient to do OAM only in the presence of the traffic then hybrid is needed. We cannot monitor the protection path too - I would clarify this question to the group. It is not just OAM but a list of active OAM. I am using the terminology of RFC7999 - specially constructed packets being injected into network for management and monitoring.

Matthew: We need to address what is specific to NVO3. Many of the solutions in the past have been on addressing specific issues.

Greg: I can point to the work that is done by the Overlay OAM DT, one of them was requirements document. It may be a good starting point to see how much of that work is applicable to NVO3 and use it as a checklist. That document may not be progressed as RFC but may be a good guide what we need to do.

Sam: We do not disagree. There are proposals and they need to be progressed.

2. WG status update.

WG Chairs.

20 mins.

3. Geneve: Generic Network Virtualization Encapsulation

draft-ietf-nvo3-geneve

Ilango Ganga

10 mins.

Ilango presenting remotely.

[presentation]

Ilango: Giving an update on behalf of Geneve authors.

Tal Mizrahi/Marvell: You mentioned alternate marking, do you believe it should be supported by using the extensibility mechanism?

Ilango: We would propose using the extension for this.

Tal: Adding Geneve option for just two bits is costly. It would be better to allocate in the common header.

Ilango: There are multiple proposals, our opinion for this OAM use case is on extension header.

Greg: I agree with Tal. Using extension for two bits is too much. We already have O bit in the header and that reasonable to discuss the use of O bit and then it would need only one more additional but for the solution. I think it is reasonable.

Kyle Larose/Sandvine: The option is usable for either extensible OAM options. It is not the only use.

Greg: I would like to see other proposals then. I do not see what the O bit is for.

4. Geneve Protocol Security Requirements

draft-mgmt-nvo3-geneve-security-requirements

Daniel Migault

5 mins.

Sami presenting on behalf of Daniel.

[presentation]

Mathew: Please show of hands of who have read the latest version. A few. This needs more review for being adopted.

Alia: This is to remind to really pay attention to security. Please please please actually review the security work. We wish to have Geneve work to progress, we need to be sure that we have ways to accommodate the increased requirements on confidentiality and integrity to be covered by Geneve.

5. Applicability of EVPN to NVO3 Networks

draft-rabadan-nvo3-evpn-applicability

Jorge Rabadan

10 mins.

Jorge presenting.

Jorge: This is not a STD track document, we are not proposing this as a single solution. We have a set of assumptions in this document too. It is a clarification document why EVPN could be a suitable control plane option.

Please review the draft.

[discussion]

Jeff Tantsura: It is extremely well written document, very useful.

Matthew: Show of hands of who has read this draft? Would be good to progress it fairly quickly. We will look at issuing adoption call shortly.

6. Performance Measurement (PM) with Alternate Marking in NVO3

draft-fmm-nvo3-pm-alt-mark

Giuseppe Fioccola

10 mins.

Giuseppe presenting.

[presentation]

[discussion]

Greg: I am a coauthor. We do have documents in other WGs that demonstrate applicability of AMM to different layers as well. It is interesting for NVO3 to discuss this document. Going back to allocation of bits - it would be good to have this discussion on the list. In my opinion there are different options for this. Allocation of O and C bit options could be explained better. Experience with SFC WG it was explained that C bit is obsolete and was returned to reserved pool. If we want to be conservative how bits in Geneve header are used it would be good to clarify how currently allocated bits are used.

Ilango: It clarifies the use of C bit - there are implementations, and it is used to indicate whether that is processed in fast or slow path. It allows for separating data and OAM packets to be sent to data and exception queues, and O bit allows for that. This allows to expedite OAM packet not to sit behind data packets.

Greg: I understand that SFC WG had its own reasons to deprecate C bit. I believe there is something to be looked at to analyze the prioritization. The whole idea of OAM is that you are fate sharing with the

data that you are measuring or monitoring. Doing something special for OAM skews the results. I strongly disagree with this approach.

Ilango: We do not see any noticeable difference. The proposal of using bits - sequence numbers, additional functionality - it allows to identify specific packets for timestamping as example. It is for specific use cases.

Greg: We are not asking for timestamping sequence number. We are requesting two bits, not timestamps.

Greg: We are asking WG to discuss and hope that editors of the document will address that.

Matthew: Please take this to the list.

7. Geneve encapsulation for In-situ OAM Data

draft-brockners-nvo3-ioam-geneve

Frank Brockners

5 mins.

Tal presenting

[presentation]

[discussion]

Tal: We would like to receive feedback whether WG is interested in working on this, and also answers on the questions raised.

Greg: You state that the benefit of IOAM is to use data to export data to export telemetry data rather than use gRPC, Kafka, etc. Telemetry is some information that is state of the node, physical resources,

that is not measured per flow. That is usually exported out of band. You are suggesting to export this via live data traffic and mentioned that this is an advantage.

Tal: We have static information like node ids and dynamic information like queue depths.

Greg: You do not have NVE in the middle, what are you tracing?

Tal: Those nodes do not need to be Geneve nodes.

Greg: Those nodes are not Geneve nodes, they are IP nodes.

Tal: We want intermediate nodes, regardless they are Geneve or not, to be able to export the flow information.

Greg: You ask for layer violation.

Tal: I do not want to use that term. We want to be able for intermediate nodes to be able to update the IOAM metadata.

Sam: Please take to the list, this is a good discussion.

Greg: This is a proposal by OAM DT team and it was called for adoption earlier this year. Are you aware of this?

Greg: The option that you prefer to have in-situ OAM after the Geneve header - that was exactly the proposal by the OOAM DT.

Sam: What is the question?

Greg: Whether you are aware of OOAM DT work?

Tal: Is that good or bad?

Greg: Not asking whether it is good or bad, asking whether you are aware.

Uma Chunduri: Is there a max limit on the OAM data that can be located in the packet - 300 bytes, 500 bytes, or as much as the MTU supports?

Tal: There is max length field in the header, It defines the length for the trace option.

Uma: If Geneve has limitation of 128 bytes, then there is no max limit.

Tal: That is exactly one of the points there.

Uma: What is the disadvantage of the left hand option?

Tal: You can skip all the options at once without interpreting them.

Frank Brockners: To be clear - left side does not exist today. Geneve has its own protocol type. Tunnel option is the only option which is a challenge with implementations. This leaves us to 128 bytes max and that is a concern. We could spare a bit or two for length in Geneve header then encapsulation should be more workable.

Frank: There is a notion of transit hop. There is a notion of end to end. Sequencing, timestamps, it depends on the use case how you use those things. What we want is a container to carry the IOAM data.

Sam: Cutting the line.

Matthew: Discussions to the list please.

8. MAC move/flush over Geneve encapsulation

draft-boutros-nvo3-mac-move-over-geneve

Jerome Catrouillet

5 mins.

Jerome presenting.

[discussion]

Ali Sajassi: A few questions/comments. This is for data plane learning. For data plane learning, how you do it for VXLAN?

Ali: There is a solution, this issue has been around for a while. To make encapsulation transport independent, RFC7769? suggests to use IEEE type. That makes it transport independent. Have you given any consideration to existing solutions?

Sami: Yes, it can be used to flush.

Ali: There are two options, both MAC withdraw and in-band.

Sami: We have MPLS-TP MAC withdrawal. We send an in-band message.

Ali: If you have solution for one, do you need to reimplement the wheel for the other?

Ali: MAC withdraw is not transport dependent. You need to invent another mechanism for GPE, for GUE.

Jerome: WG is moving to Geneve.

Ali: There are good reasons for using transport independent mechanisms.

Alia: When you talk about the ACK, this is per node and you need to keep state of which is acknowledged and not.

Jorge: You are assigning VTEP ID, but there is no information based on what. It needs much more information. Redundancy group - how does the active node know the standby? There is no way to find that.

Michael Smith/Cisco: I assume this is not piggybacking on dataplane packets?

Jerome: Correct.

9. Geneve applicability for service function chaining

draft-boutros-nvo3-geneve-applicability-for-sfc

Sami Boutros

5 mins.

Sami presenting.

[discussion]

Kyle/Sandvine: Do you discuss how you map the packet back to the path?

Sami: There are two options mentioned in the draft. One is SFF becomes stateful itself, or passed the metadata back.

Kyle: What is the benefit over those two options? Why would anyone do this instead of SRv6?

Sami: Out of time, please take to the list.

?: We have a draft where we do a similar approach.

10. BFD for VXLAN

draft-spallagatti-bfd-vxlan

Greg Mirsky

5 mins.

Greg presenting.

[discussion]

Greg: Request to working group - please discuss this document. There are implementations. Should GPE be a different document or it should go to the same one?

Alia: I am happy to see VXLAN work. GPE is not going to be STD track, and WG needs to be focusing on OAM for Geneve.

11. Group Policy Encoding with VXLAN-GPE

draft-lemon-vxlan-gpe-gbp

John Lemon

5 mins.

John presenting.

[presentation]

John: Eventually we would like a codepoint for this.

[discussion]

Matthew: Please take any comments of this draft on to the list. The focus in the WG is on Geneve.