

Minutes for the DetNet 112 WG Session

Please help capture the discussion in-line below.
No need to cover what is on the slides, just the discussion.
Please also (optionally) add your name here:

Note takers:

Ethan Grossman

DetNet Agenda for IETF 112 (Virtual)

WG ICS: <https://datatracker.ietf.org/meeting/upcoming.ics?filters=detnet>
WG Documents: <https://datatracker.ietf.org/wg/detnet/documents/>
Jabber: <xmpp:detnet@jabber.ietf.org?join>

Joint Session with PALS and MPLS on Monday Session III – November 8

17:00-19:00 (CET, UTC+1) – 16:00-18:00 (UTC)
Time Zone Converter: https://www.timeanddate.com/worldclock/converter.html?iso=20211108T160000&p1=1440&p2=tz_cet
Session ICS: <https://datatracker.ietf.org/meeting/112/session/29070.ics>
Meetecho: <https://meetings.conf.meetecho.com/ietf112/?group=pals&short=&item=1>
Audio stream: <http://mp3.conf.meetecho.com/ietf112/pals/1.m3u>
Materials: <https://datatracker.ietf.org/meeting/112/session/pals>
Note taking: <https://codimd.ietf.org/notes-ietf-112-pals>
Jabber: <xmpp:pals.ietf.org?join>
Jabber Logs: <http://jabber.ietf.org/logs/pals>
WG Documents: <https://datatracker.ietf.org/wg/pals/documents/>

Wednesday Session I – November 10

13:00-15:00 (CET, UTC+1) – 12:00-14:00 (UTC)
Time Zone Converter: https://www.timeanddate.com/worldclock/converter.html?iso=20211110T120000&p1=1440&p2=tz_cet
Session ICS: <https://datatracker.ietf.org/meeting/112/session/29068.ics>
Meetecho: <https://meetings.conf.meetecho.com/ietf112/?group=detnet&short=&item=1>
Audio stream: <http://mp3.conf.meetecho.com/ietf112/detnet/1.m3u>
Materials: <https://datatracker.ietf.org/meeting/112/session/detnet>
Note taking: <https://codimd.ietf.org/notes-ietf-112-detnet>
Jabber: <xmpp:detnet@jabber.ietf.org?join>
Jabber Logs: <http://jabber.ietf.org/logs/detnet>

Available post meeting:
YouTube:

Slot# Start Duration Information

#1 13:00 10 min Title: Intro, WG Status, Document Status

Presenter: Chairs

Status: see slides
Liason received from ITU-T SG13 Q6, response will be discussed on list.

#2 13:08 10 min Title: OAM Framework

Presenter: Greg Mirsky
Draft: <https://datatracker.ietf.org/doc/draft-ietf-detnet-oam-framework/>

Lou Berger: The Metrics section seems light, I expect that you plan to improve it. Also I think it would be good to have requirements for each metric discussed, and "must, should, may" statements for each.

Lou Berger: (Note that no objections were raised to the changes proposed on slide 6)

Lou Berger: Regarding hybrid discussion, off-path/out-of-band. In TE we have done a lot of work in support of out-of-band for reverse/bidirectional OAM. That definition is a bit different than RFC 7799 - they say "active and passive", and "passive" is "out-of-band".
Greg Mirsky: RFC7799 has 2 types of hybrid. The title of the RFC is "Active, passive and in between" so their hybrid methods that add info to data packet are hybrid, since they mark packets to identify changes in batches.
Info collected at the node (out-of-band collection methods) are interpreted here as out-of-band with respect to monitored data follow.
Lou Berger: out-of-band for responses is well established, so we want to support that. The text in the doc is good, it is just the characterization here that I was questioning.
I think it would be good to better describe what was meant by "Telemetry" in the draft.
Greg Mirsky: Telemetry here means "transporting session information". Protocols used can be characterized as on-path. Others are model-driven based on yang push.
Lou Berger: So it is more about getting it to the central controller rather than about on-path data collection? That is confusing to me.
Greg Mirsky: I didn't go into the details about discussions in IPPM WG on "iOAM direct export". OAM data is not collected in packet but is generated, then the info processing is determined by local policy so could be collected by ingress or egress node, or a controller e.g. for analytics.
Lou Berger: It would be good to bring that discussion into the document to make clear how you are using telemetry.
Greg Mirsky: OK.

#3 13:25 10 min Title: OAM Functions for The Service Sub-Layer

Presenter: Balázs Varga
Draft: <https://datatracker.ietf.org/doc/draft-varga-detnet-service-sub-layer-oam/>

Lou Berger: Regarding moving text to the existing working group document - why is this a separate doc? It would be better to get agreement from WG to merge, then merge as soon as possible.

#4 13:34 10 min Title: MPLS OAM

Presenter: Greg Mirsky
Draft: <https://datatracker.ietf.org/doc/draft-ietf-detnet-mpls-oam/>

Greg Mirsky: New d-ACH label format was presented to joint PALS/MPLS/DETNET session - generally positively received.
Lou: It would be good to hear from the WG any objections or comments about bringing this text into the WG draft.

Fan Yang: What is the procedure between this proposed new d-ACH format and Open Design Team result? Do we import this then align with DT? What is the process of how to integrate these?
Stewart Bryant: Regarding ownership of structure of nibble 1 of ACH, that probably belongs in PALS where it was created. But there has been no proposed use for those bits so far in the last 20 years. So would be good to socialize in PALS/MPLS space to make sure nobody can think of long term issues with it. But making it a parameter of the channel type is probably OK; that's what you are proposing, right.
Lou Berger: We are coordinating with that group, and we need to continue with that. We need to document what is our WG position; so far we only have an individual proposal. So if this proposal is OK with the WG then we should bring the text into the WG draft and socialize it - eventually it may even be moved to different draft. But we need to be clear on individual as opposed to WG position.
Stewart Bryant: Good to have that discussion sooner rather than later.
Lou Berger: Wasn't this already presented in the joint session?
Stewart Bryant: I don't recall that.
Greg Mirsky: I presented the DetNet ACH format at one of those meetings.
Stewart: OK.
Lou Berger: In parallel with WG, we must socialize and coordinate with the joint design team and be aware that anything we do here could be affected/modified, and/or may end up in, the PALS WG. But it is important to put together a position that is good for DetNet, then coordinate across WGs to find solution that is good for all.
Stewart Bryant: Need to see if anyone in PALS might think of an issue with this; probably not but we need to check.
Lou Berger: Then we can have a joint last call.

Greg Mirsky: I feel that this topic is important, and would like to see a conclusion from WG on updating the WG draft with this text. Regarding Design Team work, they are creating registry for first nibble values - value 1 is reserved for PW-ACH or PW-DetNet-ACH. Any new mechanism should not impede existing functionality so probably this would use a different first nibble value. We can continue this discussion on the WG list.
Lou Berger: As editor of WG doc it would be good if you could conduct a poll on bringing in the text from Balazs' doc.
Greg Mirsky: OK.

#5 13:49 5 min Title: PREOF for DetNet IP

Presenter: Balázs Varga
Draft: <https://datatracker.ietf.org/doc/draft-varga-detnet-ip-preof/>

Toerless Eckert: Need to discuss the issue of how elimination function (the jitter and buffering that implies) can be integrated into the bounded latency calculations so as maintain accuracy of calculations for queueing etc.

Balazs Varga: I thought you were discussing POF and not this draft?

Toerless: No, it is PEF. That is where it creates differences in latency. Differential path latency between A and B paths. Would be good to drop into queue without causing jitter. May need some buffering for this to simplify bounded latency calculations

Lou Berger: What does that have to do with this doc? Isn't this a general statement about PREOF? Do we currently have an actual document regarding this topic? The topic is more generally applicable to an existing RFC and any data plane, i.e. about bringing techniques from MPLS to IP. So it is a valid comment but not about this document, since this is about adopting an existing technique, but your comment is about that technique, so not specific to this draft topic.

Toerless Eckert: Agree, it is more general.

Lou Berger: So let's have a more general discussion, not just in context of this doc.

Stewart Bryant: Why do we need to have this document? We have the MPLS doc and how to do MPLS over UDP/IP. Why doesn't that just work by inheritance?

Balazs Varga: The main intention to make it clear and obvious how to do it.

Stewart Bryant: Maybe we just need a single line added to an existing draft saying that "you can use MPLS PREOF"? But maybe not worth a whole doc?

Lou Berger: We have done similar things in other groups (i.e. have a separate draft).

Pascal Thubert: With encapsulation, for IP, you lose the ID of the inner flow unless you do deep inspection. So what identifies a copy of the packet, like sequence number and flow ID is typical, but here we are tunnelling with nothing added to headers, so question arises of how to find flow?

Balazs Varga: The sequence counters are only needed on relay nodes...

Lou Berger: We can take that discussion to the list.

Pascal Thubert: Probably need more than just a tunnel for IP data plane.

Lou Berger: In other words what you are saying is that we should leave room for true IP sol'ns in future, right?

David Black: If MPLS over UDP/IP data plane used end-to-end then not much more needs to be said. But if encapsulation over only part of the path, not end-to-end, then there is more to be done, as Pascal was saying.

Lou Berger: There is some disagreement so we should continue this discussion on the list.

#6 14:02 5 min Title: Packet Ordering Function

Presenter: Balázs Varga
Draft: <https://datatracker.ietf.org/doc/draft-varga-detnet-pof/>

Toerless Eckert: We can't take the bounded latency calculus out of this discussion, otherwise it is just hand waving, pushing that consideration from this document to some undefined later document. We need to include here for this function how to calculate the latency, to make it linear to some other component. We can discuss later on list.

Lou Berger: It would be good to have an implementation doc on how to manage this, maybe one of you could create this. It is a good topic, which wouldn't change dataplane behavior, but would be a necessary calculation, e.g. for controllers.

Toerless Eckert: At least say POF needs per flow shaper. Worst case would be to have shapers for both.

Lou Berger: I look forward to reading the document.

#7 14:08 10 min Title: Requirements of large scale deterministic network

Presenter: Peng Liu
Draft: <https://datatracker.ietf.org/doc/draft-liu-detnet-large-scale-requirements/>

Lou Berger: Regarding slide 10: Requirements are good but in a presentation like this you don't need to do comparisons or gap analysis, since this will cause discussion and possible argument, so this diminishes the document. Please continue to next slide for this presentation since time is short.

Janos Farkas: What is the relation of this doc compared to the one you presented at the interim, since this one is shown as "00"? Replacement? New?

Yizhou Li: The one presented at interim was not a mature draft - this is the first "mature" version of the draft. That was a collection of info from slides, not really a draft.

#8 14:22 10 min Title: The Requirements for Wide-area IP Deterministic Networking

Presenter: Quan Xiong
Draft: <https://datatracker.ietf.org/doc/html/draft-xiong-detnet-wide-area-ip-requirements>

Peng Liu: I think it is good to characterize these requirements, but in the previous draft I just presented we also mention these requirements but don't provide details. Would it be good to incorporate this info into the previous draft?

Quan Xiong: Documentation should provide such details. We can do the research to provide details of classification of DetNet services, not just latency and jitter, but also the other requirements in your draft.
Lou Berger: Good to discuss on the list how to merge the documents and/or find places of agreement/disagreement.

#9 14:30 10 min Title: Micro-burst Decreasing in Layer3 Network for Low-Latency Traffic

Presenter: Zongpeng Du
Draft: <https://datatracker.ietf.org/doc/draft-du-detnet-layer3-low-latency/>

(No comments)

Lou Berger: Anyone interested please contact the authors on the WG list.

#10 14:40 10 min Title: DetNet Data Plane - MPLS TC Tagging for Cyclic Queuing and Forwarding

Presenter: Toerless Eckert
Draft: <https://datatracker.ietf.org/doc/draft-eckert-detnet-mpls-tc-tcqf/>

Janos Farkas: Regarding WG adoption: We are developing requirements, which goes before solutions. So wouldn't it be better to finish requirements first?

Toerless Eckert: We have been doing this for 4 years, and we know the industry needs it. Any such requirements docs would have wider scope, but just addressing a bounded latency solution should not need to wait for all of that - it could be done in parallel. Some WGs even prefer to develop a proposed solution in parallel with the requirements.

Lou Berger: Chairs can discuss any appropriate changes in chartering with the ADs, then bring to WG for comments. It would be good to have more discussion around this doc and the requirements, in parallel.

Toerless Eckert: Those requirements drafts mostly address "novel things" as opposed to short term things like PREOF.

Lou Berger: We have discussed before that queueing is not usually done in the routing area - so we need to make sure IESG and AD would be on board with us doing this. In the meantime we can run in parallel and try to have mature text by next meeting.

Toerless Eckert: In one joint meeting we discussed that work on PREOF is also QoS function?

Lou Berger: Still need to work with ADs on charter for this, that discussion doesn't change that.

From the chat window:

János Farkas: Note also: <https://www.ieee802.org/1/files/pub...9/df-finn-multiple-CQF-0919-v01.pdf05:46:54>

Shaofu Peng: It is still not clear for the congestion, for an intermediate node, multiple incoming ports from many sources receive packets in the same cycle. These packets are put into the same cycle-related queue. Can they be guaranteed to be sent together in the next cycle? It should provide more details to describe that. So far, the examples of t-cqf I have seen are all single flows to describe their processes, and the description of multiple flows is very necessary.

#11 14:50 10 min Title: BGP Flow Specification for DetNet Flow Mapping

Presenter: Quan Xiong
Draft: <https://datatracker.ietf.org/doc/html/draft-xiong-idr-detnet-flow-mapping-00>

Jeffrey Haas: The Base flowspec RFC is known to not be extensible so we are working on a v2 of that in IDR. Some of your requested changes will require flowspec v2 to become viable. that work should be ready by early December, and we would invite you to present there.

Quan Xiong: I will change reference to flowspec to v2. But for this preso would like to get some feedback on flow mapping requirements.

Lou Berger: I think you need to clarify that this document is focused on TSN mapping, not generic DetNet, so it would be good to specify it explicitly as TSN in the description. Regarding process, DetNet chairs are OK to work on this topic in either DetNet or IDR (where flowspec is defined).

Jeffrey Haas: We have time in IDR tomorrow if you want to present there.

Lou: If that happens please then let us know the opinion of the IDR group - that would help us find where to put this.

Jeffrey Haas: Also this will help DetNet determine whether this makes sense for DetNet.

Lou Berger: Thank you all, perhaps might be in person at next meeting. Bye.

Adjourn 15:00