

Start time: 3:23pm

Chair in person: Bill Cerveney

- Note Well
- Sarah Banks note taker
- Jeferson Nobre – Jabber Scribe
- Blue Sheets

- Update since IETF 90
- Calls for adoption
- Q: Greg Mirsky - Clarification on draft submitted to IESG (without closing comments?)
- A: Bill Cerveney: We declared rough consensus, Al will explain the changes with the -07 draft.
- Q: Greg Mirsky – We agreed to make changes in Toronto, the changes were made, but they made the document inconsistent
- Comment: Al Morton – That’s your opinion, but we’ll have a presentation this meeting on that, we can discuss.

- WGLC for ippm-ipsec

- Al: Some concerns over applying this to owamp. (general discussion)

Agenda:

1. Intro/Status/Adoption/Agenda
2. Working Group Documents
3. rate-problem
4. 2679-bis and 2680-bis
5. YANG model for TWAMP
6. Individual Drafts
 - a. elkins-pdm-options
 - b. Mornuley-ippm-initial-registery
 - c. Morton-active-passive

Al Morton Presenting – Access rate measurements on production networks

- Draft covers control protocol
- Section 5 compromise reviews (Toronto resolved a WGLC comment, reviews to change from MUST to RECOMMENDED – further amended now such that architectures are RECOMMENDED for both sym/asym sizes).
- When does asymmetric packet size control get indicated for network testing? (a long list was derived)
- Al Morton’s belief is to live with the Toronto compromise

- Comment from Michael Klobardans – I like this, there's a need for this, packet ensembles, chirps, are they well known terms?
- Al Morton: They're well known in the world yes, reference in the language of research papers, these terms are defined there.
- Comment from Richard Scheffenegger: enterprise solutions where you have asymmetrical routing, where we need the asymmetrical functionality.

2670/2680-bis Presentation

Presenter: Al Morton

Overview – Both of these RFCs can be modified based on our experience, and assuming we reach consensus.

All original co-authors are onboard – no IPR claimed

Summary of changes (slide)

Comments on the drafts – Al has addressed all comments on list prior to IETF-90 – some forgotten comments from the list, some off-list comments, they're all resolved in this current draft. See slide for the 5 changes addressed

Latest revisions from Joachim and Rudiger – wanted to pick up the concept of RFC Formed packets – the reference is in there, take a host timestamp as late as possible – minor document cleanup.

Issues to discuss further – haven't done anything with sections that have comments embedded in the section (Al thinks we should delete this sentence, where it's obsolete, or where we can reference an RFC that covers that type of testing).

Where is the line? RFC 6410 – we're at 2 levels of maturity. Discussion.

Next steps – further review? WGLC?

- Comment from Greg Mirsky, Ericsson – regarding clock sync, maybe TICTOC is an appropriate group.
- Al Morton (in response) – What they're working on isn't covered by their charter. They often come here to IPPM.
- Greg/Al discuss RFC1588 and whether or not clock sync is IP, and can be covered by TICTOC
- Al Morton asks for WGLC. Bill will put out an announcement on the list.

IPv6 PDM Destination Option Presentation

Presenter: Nalini & Mike Ackerman

Proposal – Something that’s in the basic IP transport, specific to v6, undisturbed by middle systems.

Solution – Implementation of existing extension header – DOH header (no IANA option number yet, but called the PDM option)

They (authors) want to get an end-to-end time from the client to server, get what is the response time from the user perspective, and break it into an application time, and a network time.

Nalini walks through an example, explaining how you determine the end-to-end, the server, and the network delay. Use a free tool like a packet analyzer to view the packets and review the info.

Wants to have this fit within 16 bytes, but wants granularity too.

Questions:

Q. Richard Scheffenegger – The info stored in host data is indexed by 5 tuple?

A. Nalini Elkins– yes

Q. Richard Scheffenegger – A question around granularity and other scenarios applicable to this? Can this become more complex?

A. Nalini Elkins – This was a simple example – we’ve developed 12 session types, etc, But they want to walk before they run level set, but she believes this will work in much more complex scenarios as well.

Q. Richard Scheffenegger: NFS over UDP or TCP – What is the response time? How do I determine which time becomes the response time embedded into the packet. You have to be careful when you want to do the application layer timing. Is this more network troubleshooting (at best?)

A. Nalini Elkins– Very good point, can draw out the flow, and see what happens. Perhaps it appears too simple to work?

A. Richard Scheffenegger – too much overhead for what you can gain from it. Most of this overhead you can get rid of, you can get the info from application-layer stuff

A. Mike Ackerman (follow up comment) – Did this purposely at the IP layer. They wanted it application layer independent. Maybe a draw back is that it’s not application aware...

Q. Al Morton – Richard’s point comes from the fact that you’re going to be storing state in the sender and the receiver, and there are times when you don’t need to communicate the things you’re storing. If you know your 10am timestamp, you

don't need to communicate this – but there are other timestamps which could be reduced to one.. there could be optimization in the protocol.

Comment: Greg Mirsky, discussion around the timestamps of sending time of packet 2 minus receive time of packet 1

Q: Sarah Banks – Help me understand; what the problem you're trying to solve is, here?

A: Nalini Elkins– This comes from enterprise networks – despite the fact that they can afford millions of dollars of equipment, this is intended to be a high level, first level triage tool.

Q: Richard Scheffenegger – Complexity overhead – how much memory, compute, highly loaded web-scale server, etc?

A. Nalini Elkins – No. They want the cooperation of the end systems, in the Operating Systems, that they might be able to do this in hardware. They've started the conversation with 1 or 2 OS vendors. Scalability is an issue – are these the right metrics? That's another issue – what happens in more complex situations? Another issue.

Q. Barbara Stark – This is one of the header extensions that gets dropped – 10-17% of this gets dropped

A. Nalini Elkins – I don't want to have that discussion in here – it's happening there.

Jabber question from Matt Mathis: Does this work for a back channel?

A: Don't know what this means, can Matt expand?

Q Bill Cerveny: Is this meant for the enterprise, or wide open Internet?

A. Nalini Elkins – When we started, it's for enterprises, but the gaming community thought this would be “way cool” – she sees use for this knowledge in many other applications.

Comment from Bill: I think this'll be easier to control in an enterprise network rather than over the open Internet.

Comment from Nalini: Agreed, although almost all the enterprises we've spoken to want this yesterday.

Q: Sarah – Are you looking to put together a draft with a narrow problem statement and use case, and focus?

A: Nalini Elkins – Looking for guidance from the group, to see if this is desirable.

Would we like to limit this to enterprise?

Comment from Mike Kloberdans, Cablelabs – Baby steps, starting with enterprise, but give us some direction with future plans, could we get asymmetrical path info inserted into the packets at each step of the way. Doesn't want to deal with destination options – Hop-by-hop would be dropped, where does it end, this scares us ☺ The source routing group might have something to say here.

Q. Greg Mirsky – Don't know the source routing group – but maybe the service function chaining group is a place to have this discussion.

A. Nalini Elkins– Thanks!

2 people have read the draft. More people should read the draft, and provide comments. Nalini Elkins will take this back to the list.

New draft – Performance metric registry

Presenter: Al Morton

Broadband networks – Want a registry where they can call out metrics by index number, and do something with it – how can we specify with precision the metrics and methods to implement and use? Want to standardize this across implementations.

Discussion of overall registry concept

Review of the registry construction

Review of some entries

Discussion of some open questions, with Al's recommendations

Q. Nalini Elkins – Confusing to me – If DNS changes from UDP to TCP, or if IP changes from v4 to v6 – does that matter?

A. Al Morton – We should look together at the DNS example in detail.

For v6, all of the IPPM work is meant to be applicable to both. We can work together on this.

Q. Nalini Elkins – If DNS changes from TCP to UDP, what happens?

A. Al Morton – Then you get another entry in the registry.

Q. Marius Georgescu – Statistical summarizing function – mentioned that the mean was used, but most RFCs I've read say that mean was the function advised, but why not consider median?

A. Al Morton – 2679 preferred median at the start, because they wanted to have statistics that were more stable to great outliers, there's some discussion in that. But what we realized in practice, it was representing wide spread use, most people talked about the mean of the packets that arrive in the timeout of the packet – if the packets arrive, this is the mean of the packet.

Q. Jason Weil – New metrics getting into the registry – going through IANA – a new test requires a new metric in order to get into the registry?

A. Al – Not necessarily – but there needs to be some reference to the metric, someone needs to be able to review this in some form.

Q. Jason Weil – Do you expect this draft to preload some metrics into the registry?

A. Al Morton – Yes, we expect this to be a feeder list.

Q. Barbara – Does the draft ack other orgs to use the registry?

A. Al Morton - The draft does, the registry draft does. If it doesn't already, it should.

Q. Marius Georgescu – What happens to metrics that are using other metrics?

A. Al Morton – A sort of derived metric? We have this classification of metric in IPPM, we already understand what they are

4 people have read this draft

Bill Cervený: Next steps?

Al Morton: No one has fed back comments on the list yet – Al has answers for the questions from the slides. Can the implementers implement based on this? We'd like feedback on this.

Bill Cervený: Feedback to the list please.

Active and passive metrics and methods (and everything in-between) Presentation

Presenter: Al Morton

Overview of how we got here, some attendees with history.

Passive and active have been used, but haven't been formally defined.

Greg Mirsky: this exists

Al Morton: it exists, but is new.

Philosophical debate ensues over whether or not to define active/passive in the way that they've been used in wide spread use, or whether or not to change it to some undefined thing moving forward.

Greg Mirsky invited Al to work on an existing draft.

Discussion of definitions

Next steps – Typo changes, observation points, hybrid examples discussion

Comment from Nalini Elkins – We've had lots of fights over what's active, what's passive, the framework doc could take a lot of these things, thanks.

Comment from Paul Cloverdale – Thinks this is a good document – one thing that came to mind, in other bodies, these things can be called intrusive and nonintrusive – or reference and full reference. Would it be useful to include those types of definitions into this draft?

A: Al Morton – Yes, I'll give that some thought, that's an interesting angle.

Comment from Tiziano Ionta – 1. There are more dimensions in the graphs, we have to work to reach that bar. 2. About the definitions – Have to tune or enrich the definitions. As an example, in Telecom Italia, we use the method of marking the traffic in a coloured way. This might be a passive method, there's no increase in traffic load – but on the other side, it could be active, because it has fields dedicated to management (it could be considered hybrid)

Comment from Greg Mirsky: Doesn't like to look at it this way – measurements that can modify the data flow in a non-intrusive way – the marking method doesn't alter how the packet is treated in the network, he believes this is still a passive measurement.

Al: It all depends on how the network is looking of it.

Comment from Mike Ackerman: Thanks for trying to nail this down, it'll help other teams to move ahead on their own efforts.

Send comments to the list.

Meeting end: 5:29pm