BMWG Session (IETF-110)
Thursday, March 11, 2021
Session III (Room 5)
10:00-12:00 US CST (UTC+0600) and 1600-1800 UTS
https://datatracker.ietf.org/doc/agenda-110-bmwg/

-=-
P L E A S E   R E A D T H E D R A F T S ! ! ! !
-=-

WG Status
Al: Should be in a good position to adopt new work

- Back-to-Back Frame (Update to RFC2544)
  https://tools.ietf.org/html/draft-ietf-bmwg-b2b-frame-03
  status:
    - IESG Approved in Dec 2020.
    - working with RFC Editors now

Al: Updated text provided to cover larger buffers.
Al: Any questions or comments? No questions/comments heard.

WG Drafts:

- EVPN
  https://datatracker.ietf.org/doc/html/draft-ietf-bmwg-evpntest-07
  status:
    - Returned to WG for proof reading
    - Completed Sec 1+2 proof reading plus suggestions for automated tools
    - Updated based on Sec 3+4 Review (Brian Monkman)
    - Draft was updated on Feb 16
    - Next step - Doc Shepherd - WG Last Call ends March 23rd

  - Al: Draft will be moved along after the WG LC.

- **Next Generation Firewall Benchmarking**
  https://datatracker.ietf.org/doc/html/draft-ietf-bmwg-ngfw-performance-06
  status:
    - WGLC ended in January - several extensive reviews and supportive comments
    - Draft was updated on Feb 22

- WG decision: is this an Update of RFC 3511? https://tools.ietf.org/html/rfc3511

- Brian presenting:
- Al: Has the name of the draft changed?
- Brian: Yes to reflect the wider scope of the draft.
- Al: That is a better name without “Next Generation”
- Al: Should the filename change going forward?
- Al: Author proposal is to have this draft supersede RFC 3511.
  Need to check that the WG agrees with that.
  Also need to sync that with IETF terminology.
  Obsolete definition is fairly clear.
  Do you think that means RFC 3511 to be obsoleted?
- Brian: Yes, when we look at where the technology is today,
- Al: We need to update the RFC header and introduction to reference that it obsoletes RFC 3511.
- Brian: We have done that.
- Al: The text should use ‘obsolete’ rather ‘supersedes’ if the WG agrees. Are there any objections to obsoleting RFC 3511.
- Sarah: (As a participant) Network Security Device spans a large number of devices. They are lots of other devices that are passive that are not firewalls but would be network security devices. Hence, not sure whether the title is clear enough.
- Bala: Comment in chat: No change in title.
- Bala: The same title was in -05. Wanted this to be applicable to other security devices.
- Sarah: (As a participant) I will reread the current draft, but not convinced that one draft will cover both passive and active devices. It might take a lot to be convinced. Will re-review and provide feedback. The generic term feels like a tough sell.
- Brian: Fair enough. We move xxx to an appendix, and focused on performance. We think that this draft will apply to any security device that handles traffic and makes a security decision based on that. If you agree with that then how should we change the title?
- Sarah: Didn’t mention RDS? Let me re-read it and give feedback on the list.
- Brian: Want to avoid a test lab wanting to use this draft and can use in the firewall space, but it is not meant for WAFFs?
- Sarah: Will see if I can come back with some proposals.
- Al: Didn’t hear any objections to hearing RFC 3511 obsolete. In the next draft let’s try that. That needs to be confirmed on the mailing list.
- Al: One other comment that came from my review. Rewrite of section 6.3. When I went through that, I noticed the term Throughput is still defined. BMWG has a pretty solid definition for throughput but it doesn’t for RFC 2647? Throughput is not defined in RFC 2647. RFC 2647 uses Goodput.
- Brian: A significant proportion of the security vendors did not like us using goodput.

- Al: What do you mean by throughput? Is it reliable throughput? Is it based of UDP and QUIC? Need to be sure that the vendors understand that BMWG has already claimed this term (Throughput). If we need to call this goodput and then make an explanatory statement later (e.g., device vendors sometimes call this throughput, but it takes a different definition). Can't have throughout here that is different from RFC2544?
- Brian: Perhaps we can use the wording from section …
- Sarah: Is there a reason why we cannot use the BMWG defined term.
- Brian: If we are going to take the input from BMWG back to your internal WGs, we want to have a very clear recommendation and why.
- Bala: Why we cannot use 2544? It defines frames per second, not bits per second. For firewalls we need to look at the bitrate rather than the frame rate. 2nd point: 2544 means the throughput is with 0 packet loss, and this test is slightly different.
- Al: 2544 throughput is often used at the packet layer. Need to get the right story here. I.e., don't use the same term with a different meaning. If you are saying that this is like a L4 and above throughput then that is fine, but it doesn't say that here.
- Bala: We took the same definition for 2544, the only difference is that we change the title.
- Al: That is the problem that I'm trying to fix. We can't use throughput alone with a different definition. If you talking goodput then that is okay. Can express goodput in bits per second. What to really understand what you are measuring here.
- Bala: Goodput has to eliminate retries. We need to clarify who is dropping the packets, or doing retransmission. Goodput is without any retransmission or anything.
- Carsten: I think that this is covered by RFC 2238? Need to consider who is the target audience. The question, isn't it possible to overload the term here for this layer 7 document.
- Al: If you put enough adjectives in front of it then that would work.
- Carsten?: E.g., Application throughput.
- Al: The definition for throughput is in 4 places: RFCs 1242, 1944 (correct), 2544 (missing some words), errata for 2544.
- Al: This is the definition that we use. This is the solution foundation. I think that would be a good way to go.
- Maciek: Thanks for driving this work. A good opportunity to standardize the nomenclature. xxx. Want to care about transaction throughput. Excellent opportunity at defining this term well so that there is no ambiguity. Will be huge dependencies on the packet sizes and what protocols are being used. There will always be some other adjectives associated with what the throughput, and this is the right place to do it here.
- Al: That is a good summary.
- Carsten: Two differences to throughput definition. "Allowed" and "correct destination interface". Traffic must be transmitted on the right egress port. Would you suggest referencing a definition of throughput and refine it here. Or should we define something different.

- Al: We have got the correct definition in the old goodput definition. You have to decide at what layer this definition is going to exist at. The point of allowed/permitted traffic that picks up the firewall function.
- Maciek: As Bala or Brian expressed, we have quite complex packet processing functions on those devices. The presence of not of CVEs? We are measuring performance (including when under attack). What measurement actually matters? We need to capture that the definition filters good vs bad. Not only measuring capability. Also measuring response to malicious traffic.
- Al: Some aspects of this definition are unique. But we want to differentiate between the existing definition.
- Brian: Wondering if it would be useful to have a call with folks from the other parties so that we call agree.
- Al: Not sure that we need an interim meeting. As long as everyone from BMWG is allowed to attend.
- Brian: No limitation, but the smaller the number, the better.
- Al: Does anyone have any objection to having a call to make progress?
- Al: Any preliminary agreement will be reported back to the WG. Nothing is going to happen here without the full information.
- Rob: Really good to have a name that doesn't redefine an existing meaning of the term. Agree that having a meeting seems like a good way forward.

End of topic discussion. Now covering other changes to the doc:

- Brian: All input that we have received is great, and will end up with a much more solid document.
- Al: Glad the netsec community is open the wider review here. Next steps we have got Sarah's review and it would be good to get some additional review. Also need to nail down the obsolete part.
- Brian: Assuming Al and Maciek are interested in the call, please drop Brian a call.
- Al: Should put it in the chat.
- Brian: Should I send an invite to BMWG?
- Al: Yes, that would be good. We should be as open/transparent as possible.
- Brian: Agrees.
- Al: Let's resolve these issues and then we should be able to have a WG LC.
- Brian: Sarah, when will you be able review?
- Sarah: Will try and get a review in this weekend.
- Brian: That sounds great.
- Maciek: 1 or 2 generic comments:
  Does draft address virtualized or cloud-based offerings?
- Brian: We are looking at this yes?
- Maciek: Is this a target for the draft.
- Brian: Just starting to think about it.
- Maciek: If it does, then this draft should be acting as a foundation. There is a direct impact on sizing.

- Maciek: Comment 2: What is the recommended vs optional features. E.g., should be enabled. IPS is an example.
- Brian: Purpose of our recommendations, wasn't necessary to test the effectiveness of the features. Want to ensure that the features are enabled during performance test. Not wanting to make this document an exhaustive security efficacy test.
- Maciek: I will revisit this. Some/many of these features could be tested in isolation. Will review section 7.1.
- Brian: We welcome comments and suggestions. Concerned if we change the definition of the tests.
- Maciek: I don't expect my review comments to be changing the definitions.
- Al: Should be able to get this out to the mailing list soon.

- Multiple Loss Ratio Search
  draft-ietf-bmwg-mlrsearch-00.txt
  https://datatracker.ietf.org/doc/html/draft-ietf-bmwg-mlrsearch-00
  status:
    - WG Adoption: successful
    - Comments (many previous questions) on the list:
      https://mailarchive.ietf.org/arch/msg/bmwg/DdEqW8kT54-PNtiXFNv3FYHh8go/

  Vratko Polak presenting:

  - Al: It will be interesting to see how you handle the device performing in a non-deterministic way under increasing load. With physical devices are able to handle transient issues. For virtual devices they are more affected by transient issues. It will be interesting to see your solution here.
  - Vratko: The algorithm aims to ??? Sometimes this happens reliably. When we switch for 5 seconds to 30 seconds things can break, but they don't break as badly as before. But ???
  - Al: Think that you are heading down the right path. As a participant, this is getting better. We are a bit aggressive at moving the limits around, which only turned up in one corner case of testing.
  - Al: Any comments on the draft, or future plans?
    No comments.
  - Al: Are there any folks willing to review.
    No comments.
  - Al: We will push for reviews on the mailing list. Congratulations on your first version WG version of the document. This is where the WG really needs to start paying attention.
  - Maciek?: Sorry missed this question. What implementations are there.
  - Vratko: One implementation in python? Not aware of anyone else trying to implement this algorithm. Everyone I know is using

- Maciek: xxx-vbench guys are using the library. Al may know better. In the context of FDIO, we have members of the Linux networking foundation. Intel and Arm are two parties who are using the MLR search. We are currently testing VPP and FDIO, the other guys are testing whatever they develop. Many goals are to reduce the time to discover the rates. Want it to be part of CI/CD. Not only applied to packets per seconds, also applied to XXX. We think that it is useful for other cases.
- Al: OpenNFV has become xxx. Thank you for your discussion today.

Proposals:

- A YANG Data Model for Network Interconnect Tester Management
  https://datatracker.ietf.org/doc/html/draft-vassilev-bmwg-network-interconnect-tester-05
  status:
    - post IETF-109 comments from Raphael Vicente Rosa and Tom Petch
    - Draft updated Feb 16, 2021
    - open-source/hardware implementation of generator/analyzer at the hackathon.

  Vladimir presenting:

  - Vladimir: Next steps is to find more folks who are interested in working on this. Can continue with the next hackathon.
  - Al: Wanted to check on the WG page, one YANG warning is being reported.
  - Vladimir: Looks like it might be a bug in the tool. I will check.
  - Al: It looks like a pretty minor error that should be easy to turn that indication green.
  - Vladimir: I'm opening to get some co-authors to help with the work. Good opportunity to collaborate with the MLS search tool.
  - Vratko: MR Search uses pluggable search. I can envisage the tool being part of that measurer. MLS search seems to be a subset of what this draft is defining. Will make the FDIO code to more closely align to this. You can count on me for reviews.
  - Al: Any more comments?

- **AOB**
  Al: Email for KJ? on xxx draft. They will have sent an update to the list (the meeting was at an inconvenient time for them).

=-