

Benchmarking Methodology WG (BMWG)

Thursday, November 13, 2014
1640-1910 HST Afternoon Session III
Lehua Suite OPS bmwg

Remote Participation:

<http://www.ietf.org/meeting/91/index.html>
<http://www.ietf.org/meeting/91/remote-participation.html>

0. Agenda Bashing

Greeting from Al and Sarah, the chairs. Encouraged participants to join the mailing list and participate. No changes to the agenda.

Note Takers: Lucien Amramov and Bill Cerveny

1a. Charter and Milestones (Chairs)

Brief update from Al.

1b. WG Status (Chairs)

Brief update from Al.

1c. Resolve OPS-DIR comments on BGP Convergence Draft

Sue not present, Al standing in as proxy.
Editorial comments to be updated; new draft coming.

2. Traffic Management Benchmarking

Presenter: Ramki Krishnan

Results of WGLC, overview of all the comments received and next rev, (should come out mid-November)

<https://tools.ietf.org/html/draft-ietf-bmwg-traffic-management-00>

- Ramki: One author had to be removed, there have been issues with IETF systems to remove the author
- Comment from Scott Bradner: When removing an author, please place them in the acknowledgements section.
- Ramki: Already done, thanks.
- Al: All comments were perfectly valid; this draft is well on its way.

3. Data Center Benchmarking Proposal

Presenter: Lucien Avramov

New reviewers/participation

<http://www.ietf.org/internet-drafts/draft-dcbench-def-02.txt>

<http://www.ietf.org/internet-drafts/draft-bmwg-dcbench-methodology-03.txt>

-Lucien Avramov: I updated the drafts. We updated reachability, based on comments from Scott and the list. Buffer measurement methodology updates were also made, based on input from the list. That's where we stand currently.

4. VNF and Infrastructure Benchmarking Considerations

Presenter: Al Morton

Discussion on Scalability and the 3x3 Matrix

<https://tools.ietf.org/html/draft-morton-bmwg-virtual-net-02.txt>

-Al Morton: presents the NFV

-Al: This draft has become a deliverable from the new charter.

-Andrew: worth considering how an HA system would look like

-Ramki: the other dimension is around VM placement or VNF placement.

-Al: that is covered in the configuration section but we should make that more clear. Good comments thanks

-Ramki: is the scope individual VNFs or chaining?

-Al: question is about all VNF related work

-Greg: virtual function matters. so there would be a different procedure for broadband network gateway and RG?

-Al: same as physical counterparts, for the virtual entities.

-Andrew: these things can be measured independently or separately. It should also be measured together

-Ramki: with NFV we are looking at multiple sub-systems: storage, compute and energy are also to be taken into consideration

-Marius: I personally have tested in virtual and bare-metal. We had very different results.

-Al will refine the document.

5. Benchmarking Methodology for SDN Controller Performance

Presenter: Bhuvan Vengainathan

Revised draft, comments on the list

<http://tools.ietf.org/html/draft-bhuvan-bmwg-of-controller-benchmarking-01>

-Bhuvan presents benchmarking for SDN controller performance

-Sarah: would makes sense to split this draft in two, but it's for you to decide if you are ok with that

-Andrew: you can't ignore the performance of the devices being controlled. Anything the controller does is depending on the size of the tables. Things have to be measured depending on the size of the table (number of flows, expiry of

flows).

-Scott Bradner: how do you test for flow expiration in a black box? What do you look at to see when a flow is expired?

-Marius: have you tested anything with this methodology? any controllers?

-Bhuvan: yes we have tested it, and we have published results against ODL, it's published in our website blog.

-Sarah: I wonder if it might make sense to put additional testing examples in an appendix?

-Marius: the concepts are hard to grasp, use case examples can help the reader.

-Bhuvan: we will give examples and analogies thank you

-Andrew: OOB / IB management will be different. What happens with IB management when one switch shuts down or becomes unresponsive, this should be covered?

-Bhuvan: fair enough

-Fernando Calabria: Never use ping or icmp, because there are issues in measuring the link latency.

Short Updates:

6. Software Upgrade Benchmarking document

Presenters: Sarah Banks

WG Adoption, discussion, etc.

<http://tools.ietf.org/html/draft-banks-bmwg-issu-meth-05>

-Has been adopted and working on bmwg 00 draft.

7. IPv6 Neighbor Discovery

Presenter: Bill Cerveny

Ideas to address Toronto/IETF90 discussion

<http://tools.ietf.org/html/draft-cerveny-bmwg-ipv6-nd-05>

New Proposals

8. Benchmarking Methodology for Virtualization Network Performance

Speaker Given name Family name: Rong Gu

Brief statement of objectives and issues that need to be discussed and resolved via presentation during the meeting: introducing a benchmarking methodology for virtualization network performance based on virtual switch.

URL: <http://www.ietf.org/internet-drafts/draft-huang-bmwg-virtual-network-performance-00.txt>

-Scott Bradner: what is the effect of faith sharing? if the testing is overloading something, it will overload all three so how can you tell what's happening?

(Question about slide 3)

- Rong: we are doing some research on it, maybe I can come up with a better idea
- Andrew: To do time-stamping you have rely on a different device.
- Marius: why not have a model with a tester that sends and receives at the same time and there is a black box in the middle so you don't care about this.
- Lucien: needs a more simple type of testing. DUT and virtual tester; have them them on separate server. I find this too complex.
- Le Ling(?): This is a very common use case we want to measure.
- Yasuhiro Ohara: It depends on the architecture.
- Marius: If you test with physical tester, then if you compare with virtual testers, you will have the difference from of the virtual environment.
- Andrew: Cycle counters don't work. You cannot trust them. Not saying you can't do this, but you have to be really careful.
- Scott Bradner: how about testing jumbo frames? Please add these
- Scott: throughput is a loaded term it has a very specific meaning. I don't think you mean that here. I think you mean data rate here, not throughput. This is a terminology point.
- Al: (CPU/Memory measurements) White box measurements. Not benchmarking, but they can be used for reference.
- Scott: these things are not accurate and very variable (in comment of slide 6).
- Al: Recommends everyone read AQM "evaluation" draft (looks like draft-ietf-aqm-eval-guidelines-00)
- Al: You've introduced an exciting topic that we're going to get a lot of feedback on.

9. IPv6 Transition Benchmarking

Presenter: Marius Georgescu

<http://tools.ietf.org/html/draft-georgescu-ipv6-transition-tech-benchmarking-00>

- Scott: you have to translate 6 to 4 direction versus 4 to 6 so you may have difference performance.
- Scott: you are right I have not seen scalability testing, I did not come up with a factor but with a illustration.
- Andrew: seed device may need to have more than one flow each.
- Andrew: PE devices need to be able to handle a lot more devices than CE device.
- Fernando and Bhuvan have agreed to read the draft and provide feedback.
- Al: Did you measure performance with both TCP and UDP? We tend to avoid TCP measurements. Matt Mathis has said, "don't go there"
- Scott: All of these tests are not real world. You do benchmarking to have something to compare with.
- Yasuhiro Ohara: TCP and UDP are really different things. TCP might be something you want to benchmark.
- Nalini Elkins: Things that you should do is DNS lookups. Are you looking at this?

-Andrew: Measuring performance with UDP stream is more reliable than a TCP stream. Packets could "look" like TCP packets ...

LAST. AOB

=====