

Benchmarking Methodology WG (BMWG) – IETF-88

Friday, November 8, 2013

2 sessions: 11:20 – 12:20 & 12:30 – 1:30 | Plaza A | OPS | bmwg

This report is arranged in 2 parts, a summary, and detailed minutes.

This report was prepared by Sarah Banks, based on the minutes/notes taken by Bill Cerveny and Jacob Rapp. Bill Cerveny also monitored jabber.

Summary

BMWG met with 19 people in attendance in person, with 0 on the jabber. The meeting began (roughly) on time and ended 8 minutes early.

No new RFCs were published in the interim between IETF meetings.

The traffic management benchmarking and data center benchmarking drafts received extensive amounts of discussion and feedback. See below for details. The IPv6 Neighbor Discovery, Software Upgrade benchmarking and BGP Convergence drafts also received a fair amount of discussion and feedback.

Re-chartering is still top of mind for the chairs and authors with pending drafts; we are hoping to re-charter this year.

Detailed Minutes

Meeting start – 11:25am

Welcome – Al Morton

- * Whose attending for the first time? – About 6 people
- * 19 people in the room
- Agenda Bashing needed? - No bashing of the agenda

Items 1 – 3 of the agenda covered by Chairs; agenda bashing, current status of WG, re-chartering

4. Traffic Management Benchmarking

Presenter: Barry Constantine

Comments on the draft (two revs since Berlin)

<http://www.ietf.org/internet-drafts/draft-constantine-bmwg-traffic-management-02.txt>

- Slide 6 - Mukhhtiar Shaikh – Q: For each trial, bytes shaped, what is the actual pps traffic?

[A]Barry: One burst of 128K bytes back-to-back and is highlighted in the previous slide

- Slide 6 - Dean Lee – Q: Can you describe the intervals?
[A]Barry: Took a packet capture, then wrote a script looking for the back-to-back frames coming out.
- Need to clarify the actual precision for the measurements.
- Barry - Q: During the work we do, if we come up with scripting that can be open-sourced with scripting, is there a tools section:
[A]Al: Lets post it on the supplemental web site.
- Comment: Dean Lee – It’s well written and tests a lot of use cases, the question is this test seems to be a very detailed function by function test. In a high level it looks similar to conformance design and validation. Question is, when you try to do a system level benchmark, what are the functions, when combined together, what is the overall result from an end user prospective? Is design A better than design B. It might be better to narrow down to specific use cases.
[A]: Not a validation, but a benchmark. It is very ad hoc, and this is a first step to clarify the functions to test. We need to be careful not going into the weeds, but right now there is no measure and there is no way to compare.
- Dean is going to work with Barry to narrow down on use cases.

5. Data Center Benchmarking Proposal

Presenter: Lucien Avramov

New reviewers/participation

<http://tools.ietf.org/html/draft-dcbench-def-01>

<http://tools.ietf.org/html/draft-bmwg-dcbench-methodology-02>

- Mukhhtiar Shaikh Q: – We don’t see a specific call out for packet loss either in the definition or the other metrics? If we look at line-rate testing, but when we do the benchmarking, the packet loss could be at the application level?
[A]Al: It is included in the good put
[A]Lucien: We focus on the data center, in latency we focus on line rate when there is no drop.
- Mukhhtiar Shaikh Q:– Should we look beyond and see what happens when there is loss. The .02% drop doesn’t tell me if the application is affected or not.
[A]: Lucien: The Line Rate is not a loss, it is for tolerance and is explained in more detail in the draft. The purpose is to have a clear and concise result to compare.
- Ram Krishnan Q: – Redefine definition, are we redefining or clarifying. Perhaps we should clarify in the draft where we are clarifying and some are redefining.
[A]Lucien: We are both clarifying and redefining, and in the draft this is addressed

- Nalini Q: Is this for end user data centers, or ISP Data centers or both?
[A] Lucien: meaningful to data center devices so we are all talking about the same thing
- Nalini – Q: Is this done by the vendor or the end user?
[A]Lucien: It can be anyone, whoever wants to run testing.
- Nalini – Q: Test to be done in isolation? What is the correlation to isolation and real DC environment?
[A]Lucien: Getting closer to production environment and we are not fully here.
- Comment – Al – Need to discuss the idea of negative values for jitter
- Dean Lee – Q: The microburst test under and incast customers and is good for low latency customers, they want to measure the fairness and have no re ordering when buffering takes place.
[A]Lucien: Need to add the re-ordering
- Comment – Dean Lee - There are compromises to the snake test
- Comment – Al - Maybe put snake test in the appendix and define it once in an appendix and put down limitations
- Action – Dean Lee to help with the limitations of the snake test

6. IPv6 Neighbor Discovery

Presenter: Bill Cerveny

Continuing Proposal, suggested by Ron Bonica

<http://tools.ietf.org/html/draft-cerveny-bmwg-ipv6-nd-02>

- Scanning 2^{64} addresses is hard. Can be done, but a pain
- Number of addresses that can be scanned is limited by the available bw
- Ram Krishnan Q:- this is valuable, but is this substantial enough for an individual draft, or could it be combined with another draft?
- [A]Al: there's no issue with small, tight projects to be done and completed
- [A]Joel Jaeggli: it's easy to see this as a milestone, there are 2 pieces - 1 is testing the size of the ds, in the forwarding or control plane, the other is measuring the performance of the control plane, and the testing methodology falls out into 1 of these 2 lines
- Comment - Nalini - very valuable, but there are a ton of other things that would be great to see, neighbor cache, router advertisements. Al counters with asking her to write down that list and put it on the list, which she agrees to.

7. Software Upgrade benchmarking document

Presenters: Sarah Banks

Following previous WG discussions, a draft has been prepared which clarifies the details of the proposal.

<http://tools.ietf.org/html/draft-banks-bmwg-issu-meth-02>

- Current update to -02 addresses feedback from IETF-87, including the question on whether or not to address data center specific cases
- The authors made the decision to specifically exclude data center as a particular case, however, if the data center equipment can fit within the methodology specified, the authors see no problem with data center equipment meeting the goals of the draft.
- Bhavani Parise: Q – Do we need rollback as a separate line item?
- [A]Sarah – Rollback is not part of ISSU, rollback is made a decision to go back, which is the methodology for downgrade is no different that the upgrade, but lets focus on the upgrade process.

- Comment – Joel Jaeggli – What is being described as hitless transfer of control, but ISSU may not be the right terminology.

8. Basic BGP Convergence Benchmarking Methodology

Status: Cross-post WG draft version to idr and grow WG

Presenter: Dean Lee

<http://tools.ietf.org/html/draft-ietf-bmwg-bgp-basic-convergence-00>

- Mukhhtiar Shaikh: Q – Advertisement of routes and withdraw are covered?
- [A] Dean Lee – Yes
- Ram Krishnan: Q – Do you clear the tables or how to you verify that you start clean from iteration to iteration?
- [A] Dean - In order to be consistent from iteration to iteration, the BGP session is brought down completely to start clean.
- Action: Need to add a line to say the following iterations need to be as clean as the first iteration.

9. LDP Dataplane Convergence Terminology

Presenter: Bhavani Parise

<http://tools.ietf.org/html/draft-parise-bmwg-ldp-convergence-term-00>

- Sarah Banks: Q – When is the methodology document coming?
- [A] Bhavani Parise: trying to get it for London (IETF-89)
- Comment - Bhavani would like some feedback on the draft.