BMWG Meeting Notes
Thursday, July 24, 2014
1300 - 1500 Thursday Afternoon Session I
Salon A

Co-Chairs Al Morton & Sarah Banks
Minutes taken by Bill Cerveny, Mike Chen, Lucien Amarov
Jabber room monitored by Bill Cerveny

=-=

0. Agenda Bashing

- Use new Note Well (as of May 2014) - Please read BCP 79 RFC to be familiar with the policies.
- Last minute add, Mike Chen, VNF methodology

1a. New Charter and Milestones (Chairs)

1b. WG Status (Chairs)

2. SIP Device Benchmarking
   Presenter: Vijay Gurbani
   Robert Spark's Comments Fully Addressed, WGLC complete
   https://tools.ietf.org/html/draft-ietf-bmwg-sip-bench-term
   https://tools.ietf.org/html/draft-ietf-bmwg-sip-bench-meth

3. Basic BGP Convergence Benchmarking Methodology
   Status: Post WGLC and RTG_DIR comment resolution review.
   Presenter: Bhavani Parise
   http://tools.ietf.org/html/draft-ietf-bmwg-bgp-basic-convergence-02

- Comment from Ron Bonica: This may or may not be relevant because part of the charter is to never trust anything under test to say anything about performance. However, there is a new draft in IDR in which each router that relays the route puts a timestamp on it. This is something that might be useful if they ever implement it.

4. Traffic Management Benchmarking
   Presenter: Barry Constantine
   Call for Adoption and ...
   http://www.ietf.org/internet-drafts/draft-constantine-bmwg-traffic-management-04.txt

- Al: Clarify explicitly state that data shown are examples only
- Al: Call for adoption coming shortly

- Dave That: I read this draft this morning and I was really delighted. Two things: netperf-wrapper is a suite of tests against TCP and UDP that almost matches what you do here. It's been really useful for us. Why is AQM not in scope?
- Barry: It was starting to become a sprawling effort, although it was part of the original effort.
- Ram Krishnan: We didn't want to stall this working waiting for AQM. We thought that could be phase 2, once it was adopted.
- Al: Scott noticed that AQM benchmarking was being discussed in another working group. AQM benchmarking isn't in scope for the BMWG charter now, but it may be welcome in the future.
- Scott Bradner: A complete nit - should add appendices to your table of contents.

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

- By show of hands, 5 folks in the room had read the draft, with 5 new hands up for those who are going to read the draft.
- Scott Bradner: Buffer testing is the least reliable of the original BMWG tests, in my experience, as far as producing reproducible results. It's meaningless when there are other things going on in the device. In the data center, I'm not sure if buffers are shared or not. If the buffers are shared, you may not get repeatable results. If the buffers are isolated, you may get more useful information. If the buffers are compliant, you might get useful results.
- Lucien: It shouldn't matter what kind of buffering is occurring, what matters are the performance results.
- Scott Bradner: I don't think it's very common in BMWG documents to have an advice or considerations section; it might be useful in this case. If you are working with certain applications, you want smaller buffers. In other cases, buffer size doesn't matter. What matters the most is repeatibility. I would suggest adding some text along those lines.
- Q from Zaid AlBanna: Are we talking about physical or virtual appliances:
- A: Physical.
- Al: We will be focusing on virtual later in the meeting.
- Lucien: We will update our drafts and we will share our results if interested.
- Ramki: Proposed DCLC research group is also looking at application performance through better control of TCP.

6. IPv6 Neighbor Discovery
   Presenter: Bill Cerveny
   Continuing Progress
   http://tools.ietf.org/html/draft-cerveny-bmwg-ipv6-nd-05

- Scott Bradner: That last one sounds dangerously like conformance, not performance.
- Ron Bonica: In a way you've got a point, because we're not talking about how fast, except in one sense. Let's say for a moment that you're being port-scanned. During the period that you're being port-scanner, you probably want entries that you know represent real hosts on the LAN segment to have a little tiny grace period in which they can get in and refresh their entry, and not be bumped out by the port-scan. Now the question is, how long is that grace period?
- Scott: You may want to phrase it just as you described; otherwise, it is conformance: What is the behavior rather than what is the performance?
- Joel Jaeggli: There is a question about what you can measure what is happening when you're in duress and can you measure that. If you can measure it, you can talk a bit about the properties of the system under those circumstances. There are some interesting questions with respect to performance out there.
- Ron Bonica: The last two really do answer how you behave under duress. You can even put numbers around them. The answer to the first, stale time entry, is a numeric answer. The answer to the second two is a qualitative answer, not a numeric answer.
- Joel: It could be numeric. If you say, I have these two queues that I put these things into based on where the request for this piece of information actually came from. Am I doing this ND because I received a NS on the subnet or am I doing this ND because I received a packet from outside and I'm attempting to resolve the next hop? So, I could ask the question, if I put these things in two buckets, how many of the second one can I do. Under the same circumstances, how many of the first can I do ... I don't think there's a universal way to implement this, there actually may be a universal way to test it, in the sense that I can do one of these things while doing the other. It may be pass/fail or it may be quantitative, because I could be able to say I can do this many, when this is occurring. Or, if you do them at this rate, then everything stops working -
- Ron: To stay clearly within the cheaper, we don't want to ask questions that have pass/fail answers, we want questions with numeric answers.
- Joel: Those questions are actually probably answerable in a

quantitative way.
- Ron: The one thing we didn't test in this draft, you talked about the externally requested addresses and the internally requested addresses. We couldn't test what you do with the internal ones.
- Joel: Most people's implementations treat those the same way.
- Ron: This draft assumes that there is no difference between internally and externally originated ND requests.
- Zaid AlBanna: Is the speed at which the learning takes place embedded in one of these tests? Is that something that would be of interest?
- Ron: No, this would be a difficult question to answer.
- Bill: Interested in finding people to help build implementations that can confirm behavior
- Al polled the room - no responses. Al suggests taking the request to the list.

7. Software Upgrade Benchmarking document
   Presenters: Sarah Banks
   Following previous WG discussions,
   http://tools.ietf.org/html/draft-banks-bmwg-issu-meth-04

- Sarah: Rev 5 is in the works, in time for Honolulu. We received feedback from the list and from others, regarding, "do you trust your counters?". We are discussing adding a considerations section
- One person in room had read the draft, 5 others intend to read it.
- Scott: Expiration date is in the past
- Sarah: Duly noted

8. VNF and Infrastructure Benchmarking Considerations
   Presenter: Al Morton
   Revised Draft based on WG comments
   https://tools.ietf.org/html/draft-morton-bmwg-virtual-net-01.txt

- Scott: What is reliability in this context?
- Al: It's the ability to continue to maintain the function
- Zaid: Are VMs virtual services?
- Al: They are real VMs, like hypervisors
- Ram: This section is really great, diving into details. I think we can have a table for concepts would work very well
- Al: How you slice the table can work very well.
- Ram: As you compare VMs with physical appliances, you can make a number of vertical comparisons.

- Bhavani: How about config mgmt; is there any related benchmark one can do?
- Al: That would seem to be one of the things you need to benchmark
- Sarah: How granular are you thinking?
- Bhavani: I was thinking in terms of scale
- Ramki: SDN controller coverage/MSV cloud orchestration table would be beneficial to include
- Ramki: Compare virtual infrastructure vs. hardware are other verticals impacted (for example, security)
- Ramki: Proposed NFV research group contributions can be included (get to immediate problems)
- 5-6 people have read the draft; a couple more intend to read it.

9. Benchmarking for OpenFlow SDN Controller Performance
   Presenter: Bhuvan Vengainathan
   New draft since IETF 89, extensively discussed on the list
   http://tools.ietf.org/html/draft-bhuvan-bmwg-of-controller-benchmarking-00

- Ramki: Flow setup rate = what is the model you are using, the embedded model? If app is decoupled from controller vs. embedded then performance differs greatly.
- Bhuvan: Point taken (to Ramki), will be included in draft
- Ramki: Clustering in BGP - are you going to consider that as you go along?
- Bhuvan: Yes
- Ramki: Netconf as southbound protocol - specific benchmarking in this area might be valuable
- Bhuvan: Should be agnostic to protocols.
- 4 people said they would read the draft.

10. Benchmarking Methodology for Virtualization Network Performance
   Presenter: Vic Liu
   Introduction of Benchmarking Methodology for Virtualization Network Performance.
   URL: http://www.ietf.org/internet-drafts/draft-liu-bmwg-virtual-network-benchmark-00.txt

- Bhuvan: What exactly is the DUT of what you're trying to test - it's confusing.
- Vic: DUT will be routers

- Multiple commenters state that the DUT isn't clear
- Bhuvan: you need to clarify exactly what you're trying to test?
- Scott: Regarding 5.1.5 test result format, what are throughput units?
- Vic: Packets per second (millions)
- Scott: You need to clarify. The rates don't make any sense; they increase with packet byte size

11. VNF Performance Test Methodology
    Presenter: Mike Chen
    Overview of intended draft proposal, possibly in time for IETF-91

- Scott: Don't change resources during a test. Don't change two things at once.
- Draft upcoming on list.

<End of session>