

2023-02-22 Agenda and Minutes

New Time: 6AM Pacific Wednesday. 3PM CET, 2PM GMT, 1930 India time US is on Standard Time. Pacific time is UTC+0800.

- [Attendees](#)
- [Agenda](#)
- [Progress for NILE Release](#)

Attendees

[Al Morton Sridhar Rao](#) [Minh-Ngoc Tran](#) [Parth Yadav](#) [Trevor Cooper](#)

Agenda

main topics today: Comments on Containerized Benchmarking, Deepdive on eBPF XDP AF_XDP

Item	Desc	who	Notes/minutes
	Deepdive on eBPF XDP AF_XDP	Sridhar	Lots of detail on the topic, based on Sridhar's research. First step is to verify that a single L2 forwarder will work (Cilium's ability to use DPDK is in question) <add slides>

Discussion: Contribution on Containerized network benchmarking in BMWG session IETF-115	All	<p>Meet on 22 Feb 2023</p> <p>eBPF Acceleration Model Figure is not accurate - trying to cover too many scenarios. Need two figures, since eBPF can also be in user space.</p> <p>Statement that Cilium with DPDK forwarding app, and DPDK traffic gen testing is not possible. Cannot connect the parts in the testbed.</p> <p>Comment about last slide: it's mostly the Rx queue that contributes to performance,</p> <p>Sridhar will send a version of the word file with his text corrections included.</p> <p>08 Feb 2023 Invite Minh-Ngoc Tran to a future meeting, to check if our comments were adopted.</p> <p>11 Jan 2023 Review with Minh-Ngoc Tran</p> <p>Draft: https://www.ietf.org/archive/id/draft-dcn-bmwg-containerized-infra-09.html</p> <p>Slides: Considerations for Benchmarking Network Performance in Containerized Infrastructure</p> <p>https://datatracker.ietf.org/doc/draft-dcn-bmwg-containerized-infra/</p> <p>yangun@dcn.ssu.ac.kr Sridhar will contact.</p> <p>mipearlaska1307@dcn.ssu.ac.kr presented at IETF-115</p> <p>Comments: ETSI TST009 scenarios were preliminary and not tested. The draft extends these scenarios with Pod inside VM.</p> <p>Sridhar Rao Tested many similar scenarios - performance difference among with VM and without VM is negligible. Not really necessary to test. Minh needs to check this with colleagues Sridhar will share our teams results - part of Daniele's work. The reason people run VMs is for security reasons and resource isolation -WHEN Containers did not supply.</p> <p>Parameters section: Cannot use a CNI independent of the Networking model (SRIOV or user space determines this)</p> <p>Some Networking models listed are not relevant when performance is the priority. Are these even applicable when performance matters? maybe not. For example, their networking model includes Kernel space switching will perform poorly. Minh replies that ALL the Models are included. Need to categorize into good/bad perf. Minh agrees.</p> <p>Performance Impacts: Major example is Number of Cores, which is not included in the Draft. Please add this. Minh agrees.</p> <p>Section 3 clarify the use of Containers and Pods - used interchangeably now. need to clarify</p> <p>Section 3 Some generalizations about CNI are not correct - Some CNI do not use user namespace</p> <p>Figure 1 What is Container Engine? is it Container Run-time.</p> <p>4.2 Some duplication on Container Network Plugin? Agree. Can use different CNI with same network Model - no difference. The SR-IOV networking model would make a big difference with other models, but CNI not so much. The Vswitch and the Networking model are closely tied together (Models with VPP are not possible with OVS, for example).</p> <p>Section 4.3 CNI does not create its own switch functions Also, which of 5 networking models does Calico fall. Figures in Section 4.3. Normal Calico does not use eBPF Calico should fall in section 4.3.1 figure, but needs to be fixed. (remove "VSwitch") and add Kernel routing tables in the user space vswitch block.</p> <p>There was a Cilium eBPF scenario that was in the slides - this is not possible with the traffic gen and DPDK.</p>
UNH transition	All	<p>Given that our Pod terminated (time-out) is there any way to recover? Or better to start from scratch?</p> <p>Can automate the IP address assignments for all the traffic links - How? Can Config.YAML really work? or is Ansible better?</p>
Testbed - IXIA support	Sridhar / All	<p>Tim Gresham asks for resolution on this.</p> <p>If we can get it, do we want it? No good without license! So, is a license still valid or available for update?</p> <p>THEN - need to ship and install in UNH, get working there - IXIA help to do this.</p> <p>04 May 2022</p> <p>Trevor Cooper says connected and powered on, but we still might have a license issue. Need Pierre's help.</p> <p>11 May 2022 this activity seems to have stalled ...</p> <p>Pod 19 also not accessible - Dan Xu.</p>

Progress for NILE Release

(summary: items 5,6,and 7 lack the necessary automated address discovery feature, defer)

see [Nile Release Schedule](#) [Nile Release Progress page](#) M4 currently due on Dec 9, M5 due on Dec 16

1	Update OS versions https://jira.anuket.io/browse/VINEPERF-673	Tasks: 1. DPDK 2. Qemu 3. Operating Systems 4. Containers	Tasks 1, 3, and 4 have been completed/merged. Not doing 2. DONE		
2	Automate setting up eBPF-based CNIs - xdp, cilium, calico. https://jira.anuket.io/browse/VINEPERF-677	setup - xdp, cilium, calico.	Patch submitted patch, will request review.		
3	Improve the ViNePerf Build Stability https://jira.anuket.io/browse/VINEPERF-675	Starting from the build to the 3 environments 1. baremetal 2. openstack and 3. K8s	Testing in-progress: either run in OS or K8s - 2 and 3 are challenging and 3 is the priority.		
4	eBPF Metrics Collection https://jira.anuket.io/browse/VINEPERF-674	Task: Develop Tool to collect metrics from eBPF programs.	Existing tool, just running it with a script Skywalking from Apache skywalking-rover possible eBPF plugin in collectd - propose for barometer or KDDI -		
<u>Tasks below are deferred from Moselle - likely Defer Again because container networking support is poor and requires significant work-arounds.</u>					
5	Epic- VINEPERF-652 : Enhance XTesting-ViNePerf Integration	Moved to Next Release	depends on 7		
6	Task- VINEPERF-658 : Enhance framework for XTesting-K8s Usecase	Partially done (reading results from output), Deployment tool.	1 task remains		
7	Task- VINEPERF-654 : XTesting-ViNePerf Integration Enhancement - Kubernetes	Will not implement due to limitations with CNIs. Moved to Next release - if CNIs support this.	Need CNI to add flows automatically in Switches (Userspace-CNI, supports DPDK, OVS, VPP). Major impediment to integrate with X-Testing Sridhar will check with Xavier if ARP resolution is supported in Prox as a switch TBD		