

2022-12-09 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](#) [Luc Provoost](#) [Luc Provoost](#)


Agenda


4 main topics today: PROX, Internship status, Nile release, UNH Transition.


Item	Desc	who	Notes/minutes
	Special Topic: Containerizing PROX 16 Nov 2022	Luc, Sridhar, Trevor, and all	<p>Update on 09 Dec 2022</p> <p>Issue remains - haven't been able to collect the logs yet.</p> <p>Notes from 30 Nov 2022</p> <p>Trouble finding the old logs: maybe in mail from Daniele</p> <p>SR - maybe give a presentation on AF_XDP https://github.com/intel/afxdp-plugins-for-kubernetes Cilium not working with DPDK - talking with Cilium group to sort this out - maybe needs telco requests.</p> <p>Also - Korea group is comfortable with Trex, that may be all</p> <p>Futurewei - has a solution for networking, but does it work with DPDK? https://github.com/CentaurusInfra/mizar SR is exploring the solution and will connect if it makes sense.</p>

<p>Intern Update 16 Nov 2022</p>	<p>Shivank</p> <p>Notes on 09 Dec 2022</p> <p>Third topo might be possible, maybe just use single CPU core, will plan to complete topo 1 and 2 for sure.</p> <p>Need status - Shivank was working on backups and installing Ubuntu,</p> <p>Notes from 30 Nov 2022</p> <p>ACTION: Propose to Tim that we complete the Pod 12 work at end of January. Need time to transition to UNH AFTER the intern project is over.</p> <p>Status: Status Check from Shivank - needs to be completed Shivank has an exam today... Needs to take-over the work.</p> <p>First look at status - then we will review the Intern feedback form.</p> <p>Notes from 23 Nov 2022</p> <p>Tim Gresham wants to know <u>time horizon for the completion of the Intern project</u>.</p> <p>Sridhar will try UNH - Sawyer or Lincoln Lavoie can get us started. AI sent mail with requirements</p> <p>IXIA was loaned to LF, intel just hosting it. Maybe it can be moved to UNH! need to investigate.</p> <p>Shivank: Status: Exams in progress. 3 tasks need to be completed. Where are we now?</p> <ol style="list-style-type: none"> 1. Infrastructure setup: Install <ol style="list-style-type: none"> a. OS b. Software c. cloud, K8s 2. Test Setup (DUT, Testing tools) <ol style="list-style-type: none"> a. vswitch- kernel module - switching solution b. CNIs c. TGen Pods d. Forwarding Pods 3. Test Runs <ol style="list-style-type: none"> a. Run Tests b. Modify 2 and repeat. <p>Shivank's Internship page with progress: Internship 2022 - Benchmarking eBPF based solutions</p> <p>Internship 2022 - Benchmarking eBPF based solutions (pdf)</p> <p>If we can complete by end of the year - stay on Intel pod 12</p> <p>Otherwise, Start to move to UNH LaaS</p>
----------------------------------	--

UNH transition requirements	<p>All</p> <p>Notes/questions on 09 Dec 2022 https://labs.lfnetworking.org/ walk-through - and answers at the 12 Dec 2022 meeting</p> <p>Lincoln Lavoie Sawyer Bergeron Justin Choquette Sridhar Rao Al Morton</p> <p>2PTLs - created pods cannot be modified - admins do not maintain access to the bookings/hosts, can add collaborators to the bookings, owner and all collaborators will get access. SR needs to login to dashboard to appear in the list.</p> <p>Need to add Sridhar's SSH key - see above. Add Sawyer and others. Collaborators have root access.</p> <p>Sawyer and team can extend the booking dates - can do 6 months or 12 month bookings - like we had on the community labs.</p> <ul style="list-style-type: none"> • Networking questions: direct connect is possible? (can't draw it) not easily possible, proximity is not guaranteed along with specific hosts... • After booking, at 6 months duration, try working with fabric first for networking. Then try static allocation and then direct connect. Note: a few back-end changes will allow changes after booking. may be able to pull hosts out of pool. • Not very worried about the BW constraints of the fabric today, host to ToR is four 25G and two 10G link. <ul style="list-style-type: none"> ▪ Sridhar's access and key - can do. ▪ What networks have internet access? does red net have it? - currently each booking has one public network - connection to node that is NAT'ed to Internet, DHCP assignment, /24 for Pod. Other nets not Internet/public, Layer 2 only. Admins can change this. ▪ IXIA - LF owns it, could move to UNH. ▪ SW Traffic Gen - do we need to change BIOS? No. ▪ These are baremetal server bookings, right? YES you get the whole thing! ▪ Jump will be over-provisioned. ▪ how does snapshot work now? shouldn't/doesn't work. Side effect of past and needs doc fix. maybe again - but not available in 2.0: Can use cloudinit files in YAML - can send docs. Can get a server to be ready with many packages installed. add commands like a script during stand-up. ▪ Need to fill-out purpose and other fields in the Resource Request - can add collabs, also Congi field is where you paste in the YAML ▪ Ubuntu 20.04 - can be upgraded after provisioning. You own it. BMC credentials get added to the bookings and are shared with the info after. ▪ Don't change BIOS password with BMC access! ▪ Ping Sawyer and Team with problems!!! <ul style="list-style-type: none"> ▪ note 'sudo do-release-upgrade' <p>Notes from 30 Nov 2022</p> <p>SR's requirements:</p> <p>1 Node for Jump Host (50% of the HW req below should be ok)</p> <p>3 Nodes Kubernetes Cluster (1 Master and 2 Worker Nodes)</p> <p>Requirements on Each Single Node Configuration of the cluster</p> <ol style="list-style-type: none"> 1. 1x Gigabit ethernet for control-plane/internet/external-access (management) 2. 2 to 4 x SR-IOV compatible NICs for dataplane (testing) At least 10GigE 3. 2 socket x (Intel Xeon E/Gold, (or ARM) at least 22 cores, at least 2.4 Ghz, approx 50mb Cache) 4. 64GB to 128Gb Memory 5. 180GB SSD and 3TB SATA HDD storage <p>Networking between the nodes:</p> <ol style="list-style-type: none"> 1. need separate management network and dataplane/testing network 2. What does the dataplane network include? ToR switch? How many hosts share the ToR switch? 3. VLAN tags needed? How many? 4. Direct cabling between NICs of Master and one worker nodes would be ideal. <p>WHAT about COSTS? does TSC need to approve? No cost - NO Approval needed.</p> <p>ACTION Al Morton - fill-out request ASAP, then ping Lincoln and Sawyer on this.</p>
-----------------------------	--

<p>Discussion: Contribution on Containerized network benchmarking in BMWG session IETF-115</p>	<p>AI</p>	<p>09 Dec 2022</p> <p>Notes from 30 Nov 2022</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 Sirdhar will contact.</p> <p>mipearlskal307@dcn.ssu.ac.kr presented at IETF-115</p>
<p>Nile Progress</p> <p>Schedule</p>	<p>all</p>	<p>Notes on 09 Dec 2022</p> <p>Notes from 30 Nov 2022</p> <p>Re-prioritize this for December, complete dev and updates... with possible shift to next release. Prioritize results for December!</p> <p>Review Sridhar Rao input 02 Nov 2022</p> <p>Started Release Plan for ViNePerf Nile</p> <p>Sridhar has submitted 2 new patches, which will complete the dev schedule - AI must review. New SW versions and New OS versions. Adding AI as a reviewer.</p> <p>Using DPDK build with Meson and Ninja tools now. new tools. The build process is different for diff each OS now, diff folders for Fedora, etc.</p> <p>Need a discussion with Fulvio on the approach to use with building the CNIs. Sridhar will consider the questions to ask and try to exchange e-mail with Fulvio.</p> <p>Pod 12 is fine now! but we are the last user of the Intel community pods... The IXIA license has expired.</p> <p>Consider collaboration with Open Programmable Interface Project (Linux OPI) - need a member to join with us and start to identify common areas of interest. Sridhar will talk to Joe White, Dell TSC Chair.</p> <p>Node 3 in POD 13 has build issues - need to update that system. system dependencies. Build was successful in other locations, so we are interested to know where other projects are doing their builds. Maybe UNH Lab.</p>
<p>12 Oct 2022 Sridhar presented at the Tech Discuss!</p>		<p>eBPF - final slides</p> <div data-bbox="690 1239 993 1543">  </div>

eBPF topic (additional details)	Sridhar Rao	<p>Sridhar presented slides, which he used to familiarize Anuket with XDP and eBPF</p>  <p>L3AF project in LFN - life cycle management on eBPF programs</p> <p>Cillium project has done some of their own benchmarking - but uses some unknown TCP: https://docs.cilium.io/en/stable/operations/performance/benchmark/</p> <p>Not much uptake in telco - still many operations issues. The Cillium community agrees that they have not made a good case for telco usage (in conversations with Sridhar).</p>
Shivank's Intern update - Archive:		<p>19 Oct 2022</p> <p>Here are Shivank's updates for this week. https://wiki.anuket.io/display/HOME/Internship+2022+-+Benchmarking+eBPF+based+solutions+repo https://github.com/Alt-Shivam/Benchmarking-eBPF-XDP Sridhar Rao commented: target to complete the Topology-1 (refer to slides that I had sent) first as baseline results to compare with previous results.</p> <p>21 Sept 2022</p> <p>Shivank's CPU not supported by DPDK . Need this for most of our benchmarking tools</p> <p>New T-rex traffic Gen have been added and a new DPDK folder.</p> <p>Sridhar shared https://github.com/intel/afxdp-plugins-for-kubernetes it is a possible option to use Shivank will investigate.</p> <p>Mail to Casey Cain about Intern compensation and feedback.</p> <p>Shivank's report for last week: 14 Sept 2022 https://github.com/Alt-Shivam/eBPF-CNI</p> <p>Pod 18 connectivity not working for several days, impacting Shivank's work. Tim G says working... Need Pod 12 working as well. There is an active JIRA ticket and comment stream.</p> <p>Shivank testing with a VM and K8s cluster - can try out locally</p> <p>Need to review mid-term Mentoring report - Al and Sridhar completed in last half of the meeting.</p> <p>ACTION Al Morton invite Maryam meeting to discuss performance (when we agree it's the right time).</p>
<p>Post review of Daniele's work 13 Jul 2022</p> <p>closure to Daniele's work and see if there is anything worth publishing, may consider less-selective conferences if novelty is less</p>	Sridhar	<p>Some issues we dig into now</p> <p>No reply from simonartxavier@gmail.com and Luc? Al and Sridhar sent new messages.</p> <p>Once we containerize prox, and we assign multiple cores to single interface, it crashes. how can we run Prox with multiple cores in a container??</p>
Tasks and Action Plan for Shivank (Archive)		<ol style="list-style-type: none"> 1. Get access to the Testbed (Pod18 now) 2. Demonstrate how to use Prox and T-Rex on Node-5 3. Start with Baremetal and eBPF: Node-5 and Node-4 4. Kubernetes Cluster Setup - Node-1 and Node-2: Install Necessary Components.

Open Programmable Interface Project (Linux OPI) Change CPU NIC architecture, more autonomy and CPU power goes to NIC in OPI, more than smart NIC now. New Trust boundary between Autonomous NIC and CPU.	Sridhar	<p>New project, governing board formed and TSC meeting set.</p> <p>Opportunity for Benchmarking: work with us.</p> <p>Need to know planned NIC speeds, Need NIC HW to benchmark.</p> <p>20 Jul 2022 There will be a test bed discussion, UNH and Keysight collaboration - using Keysight tools. We want to join the collaboration. Maybe UNH would welcome our help to some degree. Need some time to understand how the DUT HW will be obtained and installed. Also Remote access (for us, like Intel labs).</p>
Testbed - IXIA support	Sridhar /AI	<p>2nd ACTION: Do we have support from Ancuta or other IXIA person? AI Morton sent e-mail</p> <p>ACTION: License for IXIA HW - Check with Trevor</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> <p>ACTION: AI: can you add a comment in the Jira issue.. INFRA-7 ? As a PTL you can mention the dev is stuck with nodes being inaccessible. https://jira.anuket.io/browse/INFRA-7</p>
Additional Mentoring input	Sridhar /Federica	<p>13 Jul 2022 Federica</p> <p>Past meeting: python scripts. Possible combine 1 and 3</p> <p>@Sridhar will expand on the proposal descriptions and Federica will begin to ask students about their interests.</p> <p>python scripting with K8s is possible for some students in 3rd year. Projects must have resume value, add skill set or new experience.</p> <p>Sridhar Rao created some projects.</p> <p>ACTION: See proposals & review</p> <p>https://wiki.anuket.io/display/HOME/Potential+Projects+for+Student+Volunteers</p> <div data-bbox="690 1058 1193 1560">  </div>
Intern (background info)	Sridhar /AI /Federica /Shivank	<p>XDP performance Studies for Cloud-Native NFV Use Cases</p> <p>"Maryam Tahhan joins the crew to talk XDP, AF_XDP, and fast networking"</p> <p>Operations support seems to be a big issue, and performance is currently about 0.7 of VPP and DPDK</p>
K8s on Pod 12	Sridhar	will look into it

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 and 3 have been completed/merged.			
2	Automate setting up eBPF-based CNIs - xdp, cilium, calico. https://jira.anuket.io/browse/VINEPERF-677	setup - xdp, cilium, calico.				
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				
4	eBPF Metrics Collection https://jira.anuket.io/browse/VINEPERF-674	Task: Develop Tool to collect metrics from eBPF programs.				
	<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			