Procera Networks
Virtual PacketLogic Software Performance Test

Introduction

In the wake of the telecommunications industry advancements towards virtualization, many vendors are laboring to port their solution to a virtual environment, or as the ETSI NFV ISG defines it, an Network Function Virtualization Infrastructure – NFVI.

Procera Networks participated in EANTC's inaugural Network Functions Virtualization (NFV) showcase in 2013 where the tests focused on proving fundamental requirements for Virtual Network Functions (VNFs) – instantiation and provisioning, portability and elasticity. In the second half of 2014, our focus is on prototyping service provider deployments, specifically benchmarking performance and latency for Virtual Network Function (VNF) implementations.

Many network operators believe that there is a performance trade-off that must be made when moving to Network Function Virtualization (NFV) solutions. Procera explained to us that they believe one of the fundamental premises of NFV is to unlock cost savings for operators, and to meet that expectation, the performance differential between an NFV and a purpose-built hardware platform must be minimal.

Background

Procera has implemented PacketLogic/V – an NFV version of their PacketLogic Internet Intelligence Solution for network operators. The Procera solution is a data plane solution, which traditionally relied on specific hardware configurations or hardware-assist for achieving high performance.

Procera explained that they wanted to maintain performance leadership even in an NFV environment for which reason they worked closely with Intel to optimize the performance of PacketLogic/V. With this test, Procera believed that they achieved the maximum performance possible with the physical I/O present in high performance Intel servers which is why EANTC was asked to join Procera in their lab in Sweden and to verify their performance claims.

Test Results Highlights

- **NFV implementation demonstrates bare metal performance**
- **Layer 7 full DPI throughput of 40Gbit/s**
- **Average TCP Response Time under 1ms**

Test Setup

Procera procured an Intel server for the tests. While not the beefiest server in the market, the system still had dual socket Intel XeonE5-2699v3 CPUs clocked at 2.30 Ghz and 18 cores. In addition, the server had 64 Gigabytes of memory (RAM). The system is actually the highest CPU core density available in an Intel dual socket COTS server on the market today.

On top of the physical hardware, Procera installed Debian 7.6 and Qemu 1.6.2 as well as version 15.0.4.5 of PacketLogic software.

This test was to demonstrate a "bare metal" deployment, where a service provider will run a single VNF on a server to maximize the performance of the VNF.

We used Ixia's BreakingPoint Systems as the tester in our setup and emulated real, full-stack applications in all relevant test cases.

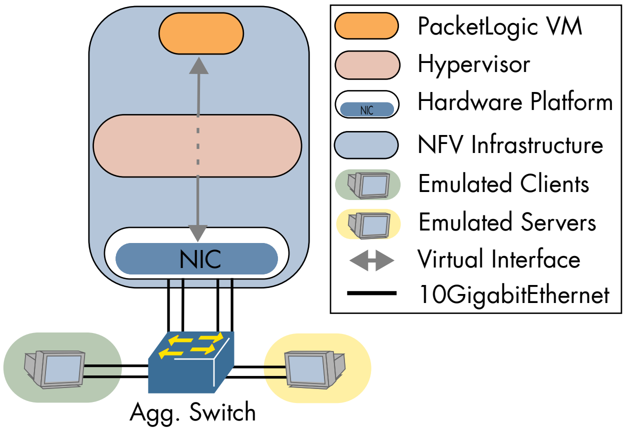

Component	Hardware / Software Details
PacketLogic Real time Enforcement (PRE)	Version 15.0.4.5
Hypervisor	KVM on Debian 7.6 with QEMU version 1.6.2
Hardware Platform	Intel Haswell server, CPU: E5-2699 v3 @ 2.30GHz RAM: 64GB
	Host operating system: Debian 7.6
	NIC: Intel 82599EB 10-Gigabit SFI/SFP+

Figure 1: Test Bed Network Topology

Results

In all test cases we had a unified goal: measure realistic performance under conditions that service providers are likely to see, and expect performance on par with specialized hardware. After all, by using a commercial off the shelf (COTS) server, service providers are looking to save on Capital Expenditures (CAPEX), while maintaining the performance of their networks. Deep Packet Inspection technology requires high performance from the underlying hardware, and any decrease in performance will immediately reflect poorly on customer experience or risk essential revenue-generating functions in a provider's network.

Layer 2 Throughput Performance

One of the baseline performance metrics of the network components (including software component) is its ability to process and forward data packets at high rate. While the function of a Deep Packet Inspection solution is clearly to inspect the application layer and its payload, in order to benchmark the performance of a system, we agreed to start with pure layer 2 switching performance.

Using this methodology we could first get an idea about the performance of the system without any of the bells or whistles and then, in the following tests, we intended to compare the performance when layer 7, full-stack, traffic is being sent by the tester. Should we then record large differences in performance, we could safely conclude that the hardware is not to blame and point our attention at the software elements. There is another benefit to using this methodology. When just focusing on the forwarding performance of the system, we could measure latency through the unit and later control for latency changes when running layer 7 traffic.

This test was run utilizing a direct PCI connection – the so called PCI-pass-through. We used standard, RFC 2899 frame sizes as well as an IMIX with 64, 128 512, 1280 and 1518 bytes packet sizes.

The results of this first baseline test can be seen in the table below. We recorded frame loss for the smallest packet sizes (64 and 128 bytes) which was unexpected but easy to understand: these test runs generated the highest packet processing demand on the system. Procera suspects this could be a limitation on the Intel IO memory management unit (IOMMU), however, neither Procera nor EANTC could confirm or deny this hypothesis at this time.

Frame size [Bytes]	Throughput		Latency [μs]	
	Tx Packet per second	Gbit/s	Max.	Avg.
64	25,000,000	16.8	6,682	152
128	19,260,000	14	4,760	152
256	18,030,000	38.9	8,823	173
512	9,398,000	40	806	228
1518	3,251,000	40	233	755
IMIX	6,689,000	40	759	273

The test showed that the pure packet processing performance of the Intel host system was indeed more than capable of performing for realistic packet sizes. As you can see, once the packet size increased to 256 Bytes, almost no loss was recorded anymore. This is actually a reasonable engineering compromise between system design and cost (power, heat etc) – Internet traffic, according to all known research (for example at www.caida.org), averages at 300-400 bytes and is certainly not all 64 bytes packets.

Layer 7 Throughput Performance

Once we had a good baseline regarding the packet processing capabilities of the solution, we moved to tests that were more in-line with the intended use of Procera's PacketLogic. In these tests we sent full-stack layer 7 traffic and measured again the throughput of the system. The throughput of the system during this test was less influenced by the packet per second performance but by actually running the deep packet inspection analysis on the flows the Ixia tester generated and creating statistics on these flows.

The first test was executed using HTTP 1.1 persistent connection for all flows. Measuring the throughput of a system using TCP is actually impossible – TCP will backoff in case packets are lost and will reduce its rate. For this reason the IETF, in RFC 2647, defined the term goodput. This term describes the number of bits per unit of time forwarded to the correct destination interface of the system under test minus any bits lost or retransmitted.

The results as recorded by the tester are shown in the figure below. As you can see, on a system with 4x10GbE interfaces, without any block or drop rules defined in the PacketLogic software, we recorded a near perfect performance.

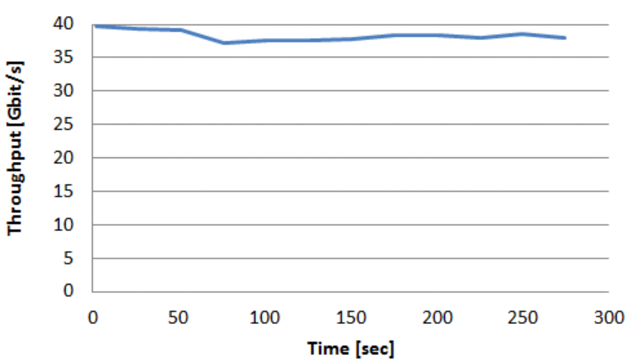

Figure 2: HTTP-Only Throughput Performance

Once we recorded this performance, we changed the tester configuration to a wider mix of traffic, one that is more indicative of modern networks and included, in addition to HTTP traffic, also YouTube and flash videos, HTTPS, some bittorrent peer-to-peer clients as well as SIP and Skype traffic. The application traffic is shown in the following figure.

Again, due to the complexity of the application breakdown we did not expect to be able to "fill the

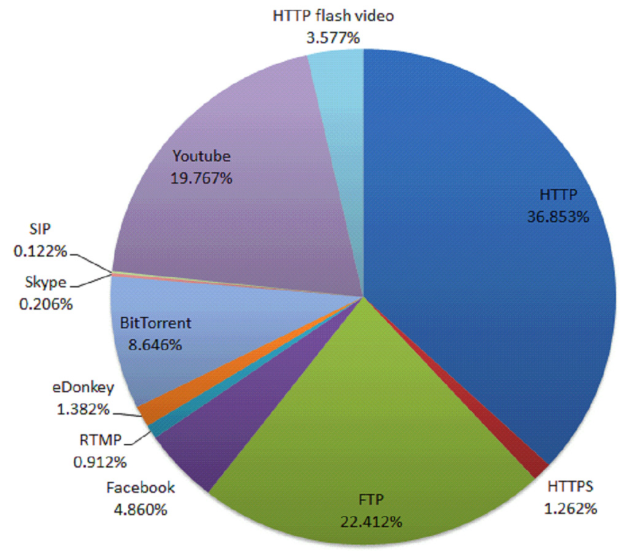

Figure 3: Traffic Profiles

pipes" to the brink (i.e. to 40Gbit/s), but wanted to see that we were able to get close to this value as well as to monitor the CPU load and control for system instabilities. As the figure below shows, the system behaved in a very predictable and reliable way and the virtual machine CPU never exceeded 45%.

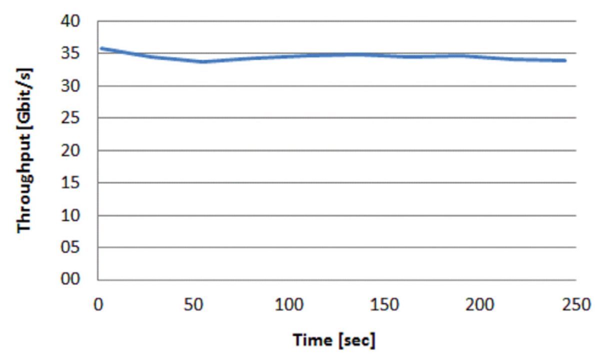

Figure 4: HTTP-Only Throughput Performance

Delay Performance

One essential key performance indicator (KPI) in a healthy network, is low delay. High delay is quickly translated, by the customer using the network, to slow network performance which leads to call-center overload and unhappy customers. Clearly, a system that is being installed in the network, with the purpose of extracting additional revenues while also providing visibility into the ways the network is being used, can not increase the delay in the network.

We used the same two conditions under which we measured the throughput performance to record delay: HTTP only traffic and our application mix. For the HTTP only traffic, Procera configured a “fast Forwarding Mode” which only inspect layer 2 and 3 while in the “Full traffic control mode” the system was configured to monitor well into the payload of the application layer.

We focused our attention on parameters that are mostly relevant to application performance – TCP setup time, response time and close time. As is seen in the figure below, when using fast forwarding mode, delay was lower than when the PacketLogic/V had to process all application layer data. The increase in latency would not have been statistically significant in a real network deployment or likely to affect network quality.

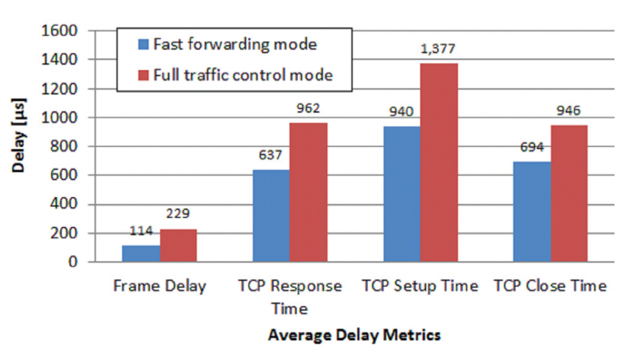

Figure 5: Delay Measurement with & without Traffic Control Function

Detection Accuracy and Reporting

Of course, the function of a traffic control system is not simply to pass traffic, but also to detect what the traffic in the wire is and to report on its nature. The goal we set for this test was to compare between the report that Procera’s PacketLogic/V provided and the report that the tester was providing, all while processing a high amount of packets.

Traffic Throttling Performance and Accuracy

The last, but certainly not least, was this classical traffic control function test. In this test, we defined certain traffic classes that the system under test had to throttle. For example, we asked Procera to configured a limitation of 64Kbit/s for all mobile subscribers (identified by their IP addresses). We also limited all peer-to-peer and Voice over IP traffic to 1Gbit/s. While these traffic control functions were being enforced we also expected that no valid traffic would be affected.

The system did exactly what we expected of it. We recorded no increase in CPU utilization when the throttling function was activated and none of the valid applications were affected.

Summary

Our goal in testing was to demonstrate that performance was not a barrier to NFV deployments for VNF data plane NFV implementations like PacketLogic. As the results presented in this document show, this goal has been met. Under the realistic service provider conditions we emulated in this test, Procera’s PacketLogic/V demonstrated performance that was only limited by the network interface card physical capacity, while still leaving enough CPU cycles free to manage the COTS server and potentially run other Virtual Network Functions on the same server.

This test also validated the vision on which ETSI’s NFV ISG is standardizing – free the function from the hardware and provide service providers with the flexibility, portability and elasticity. We hope to see new generation of servers in the near future with even more network interface cards that will then show even more performance.

About EANTC

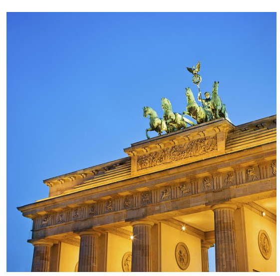

The European Advanced Networking Test Center (EANTC) offers vendor-neutral network test services for manufacturers, service providers and enterprise customers. Primary business areas include interoperability, conformance and performance testing for IP, MPLS, Mobile Backhaul, VoIP, Carrier Ethernet, Triple Play, and IP applications.

EANTC AG
Salzufer 14, 10587 Berlin, Germany
info@eantc.com, http://www.eantc.com/

V1.0, JG, 2014-10-07