       Benchmarking Virtual Switches in the Open Platform for NFV (OPNFV)

Abstract

   This memo describes the contributions of the Open Platform for NFV
   (OPNFV) project on Virtual Switch Performance (VSPERF), particularly
   in the areas of test setups and configuration parameters for the
   system under test. This project has extended the current and
   completed work of the Benchmarking Methodology Working Group in the
   IETF and references existing literature. The Benchmarking
   Methodology Working Group has traditionally conducted laboratory
   characterization of dedicated physical implementations of
   internetworking functions. Therefore, this memo describes the
   additional considerations when virtual switches are implemented on
   general-purpose hardware. The expanded tests and benchmarks are also
   influenced by the OPNFV mission to support virtualization of the
   "telco" infrastructure.

Table of Contents

1. Introduction

   The Benchmarking Methodology Working Group (BMWG) has traditionally
   conducted laboratory characterization of dedicated physical
   implementations of internetworking functions. The black-box
   benchmarks of throughput, latency, forwarding rates, and others have
   served our industry for many years. Now, Network Function
   Virtualization (NFV) has the goal of transforming how internetwork
   functions are implemented and therefore has garnered much attention.

   A virtual switch (vSwitch) is an important aspect of the NFV
   infrastructure; it provides connectivity between and among physical
   network functions and virtual network functions. As a result, there
   are many vSwitch benchmarking efforts but few specifications to guide
   the many new test design choices. This is a complex problem and an
   industry-wide work in progress. In the future, several of BMWG's
   fundamental specifications will likely be updated as more testing
   experience helps to form consensus around new methodologies, and BMWG
   should continue to collaborate with all organizations that share the
   same goal.

   This memo describes the contributions of the Open Platform for NFV
   (OPNFV) project on Virtual Switch Performance (VSPERF)
   characterization through the Danube 3.0 (fourth) release [DanubeRel]
   to the chartered work of the BMWG (with stable references to their
   test descriptions). This project has extended the current and
   completed work of the BMWG IETF and references existing literature.
   For example, the most often referenced RFC is [RFC2544] (which
   depends on [RFC1242]), so the foundation of the benchmarking work in
   OPNFV is common and strong. The recommended extensions are
   specifically in the areas of test setups and configuration parameters
   for the system under test.

   See [VSPERFhome] for more background and the OPNFV website for
   general information [OPNFV].

   The authors note that OPNFV distinguishes itself from other open
   source compute and networking projects through its emphasis on
   existing "telco" services as opposed to cloud computing. There are
   many ways in which telco requirements have different emphasis on
   performance dimensions when compared to cloud computing: support for
   and transfer of isochronous media streams is one example.

1.1. Requirements Language

   The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
   "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and
   "OPTIONAL" in this document are to be interpreted as described in
   BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all
   capitals, as shown here.

1.2. Abbreviations

   For the purposes of this document, the following abbreviations apply:

      ACK   Acknowledge
      ACPI   Advanced Configuration and Power Interface
      BIOS   Basic Input Output System
      BMWG   Benchmarking Methodology Working Group
      CPDP   Control Plane Data Plane
      CPU   Central Processing Unit
      DIMM   Dual In-line Memory Module
      DPDK   Data Plane Development Kit
      DUT   Device Under Test
      GRUB   Grand Unified Bootloader
      ID   Identification
      IMIX   Internet Mix
      IP   Internet Protocol
      IPPM   IP Performance Metrics
      LAN   Local Area Network
      LTD   Level Test Design
      NFV   Network Functions Virtualization
      NIC   Network Interface Card
      NUMA   Non-uniform Memory Access
      OPNFV   Open Platform for NFV
      OS   Operating System
      PCI   Peripheral Component Interconnect
      PDV   Packet Delay Variation
      SR/IOV   Single Root / Input Output Virtualization
      SUT   System Under Test
      TCP   Transmission Control Protocol
      TSO   TCP Segment Offload
      UDP   User Datagram Protocol
      VM   Virtual Machine
      VNF   Virtualised Network Function
      VSPERF   OPNFV vSwitch Performance Project

2. Scope

   The primary purpose and scope of the memo is to describe key aspects
   of vSwitch benchmarking, particularly in the areas of test setups and
   configuration parameters for the system under test, and extend the
   body of extensive BMWG literature and experience. Initial feedback
   indicates that many of these extensions may be applicable beyond this
   memo's current scope (to hardware switches in the NFV infrastructure
   and to virtual routers, for example). Additionally, this memo serves
   as a vehicle to include more detail and relevant commentary from BMWG
   and other open source communities under BMWG's chartered work to
   characterize the NFV infrastructure.

   The benchmarking covered in this memo should be applicable to many
   types of vSwitches and remain vSwitch agnostic to a great degree.
   There has been no attempt to track and test all features of any
   specific vSwitch implementation.

3. Benchmarking Considerations

   This section highlights some specific considerations (from [RFC8172])
   related to benchmarks for virtual switches. The OPNFV project is
   sharing its present view on these areas as they develop their
   specifications in the Level Test Design (LTD) document as defined by
   [IEEE829].

3.1. Comparison with Physical Network Functions

   To compare the performance of virtual designs and implementations
   with their physical counterparts, identical benchmarks are needed.
   BMWG has developed specifications for many physical network
   functions. The BMWG has recommended reusing existing benchmarks and
   methods in [RFC8172], and the OPNFV LTD expands on them as described
   here. A key configuration aspect for vSwitches is the number of
   parallel CPU cores required to achieve comparable performance with a
   given physical device or whether some limit of scale will be reached
   before the vSwitch can achieve the comparable performance level.

   It's unlikely that the virtual switch will be the only application
   running on the SUT, so CPU utilization, cache utilization, and memory
   footprint should also be recorded for the virtual implementations of
   internetworking functions. However, internally measured metrics such
   as these are not benchmarks; they may be useful for the audience
   (e.g., operations) to know and may also be useful if there is a
   problem encountered during testing.

Benchmark comparability between virtual and physical/hardware
implementations of equivalent functions will likely place more
detailed and exact requirements on the "testing systems" (in terms of
stream generation, algorithms to search for maximum values, and their
configurations). This is another area for standards development to
appreciate; however, this is a topic for a future document.

3.2. Continued Emphasis on Black-Box Benchmarks

External observations remain essential as the basis for benchmarks.
Internal observations with a fixed specification and interpretation
will be provided in parallel to assist the development of operations
procedures when the technology is deployed.

3.3. New Configuration Parameters

A key consideration when conducting any sort of benchmark is trying
to ensure the consistency and repeatability of test results. When
benchmarking the performance of a vSwitch, there are many factors
that can affect the consistency of results; one key factor is
matching the various hardware and software details of the SUT. This
section lists some of the many new parameters that this project
believes are critical to report in order to achieve repeatability.

It has been the goal of the project to produce repeatable results,
and a large set of the parameters believed to be critical is provided
so that the benchmarking community can better appreciate the increase
in configuration complexity inherent in this work. The parameter set
below is assumed sufficient for the infrastructure in use by the
VSPERF project to obtain repeatable results from test to test.

Hardware details (platform, processor, memory, and network)
including:

o BIOS version, release date, and any configurations that were
   modified

o Power management at all levels (ACPI sleep states, processor
   package, OS, etc.)

o CPU microcode level

o Number of enabled cores

o Number of cores used for the test

o Memory information (type and size)

   o Memory DIMM configurations (quad rank performance may not be the
      same as dual rank) in size, frequency, and slot locations

   o Number of physical NICs and their details (manufacturer, versions,
      type, and the PCI slot they are plugged into)

   o NIC interrupt configuration (and any special features in use)

   o PCI configuration parameters (payload size, early ACK option,
      etc.)

   Software details including:

   o OS RunLevel

   o OS version (for host and VNF)

   o Kernel version (for host and VNF)

   o GRUB boot parameters (for host and VNF)

   o Hypervisor details (type and version)

   o Selected vSwitch, version number, or commit ID used

   o vSwitch launch command line if it has been parameterized

   o Memory allocation to the vSwitch

   o Which NUMA node it is using and how many memory channels

   o DPDK or any other software dependency version number or commit ID
      used

   o Memory allocation to a VM - if it's from Hugepages/elsewhere

   o VM storage type - snapshot, independent persistent, independent
      non-persistent

   o Number of VMs

   o Number of virtual NICs (vNICs) - versions, type, and driver

   o Number of virtual CPUs and their core affinity on the host

   o Number of vNICs and their interrupt configurations

   o Thread affinitization for the applications (including the vSwitch
      itself) on the host

   o Details of resource isolation, such as CPUs designated for Host/
      Kernel (isolcpu) and CPUs designated for specific processes
      (taskset).

   Test traffic information:

   o Test duration

   o Number of flows

   o Traffic type - UDP, TCP, and others

   o Frame Sizes - fixed or IMIX [RFC6985] (note that with
      [IEEE802.1ac], frames may be longer than 1500 bytes and up to 2000
      bytes)

   o Deployment Scenario - defines the communications path in the SUT

3.4. Flow Classification

   Virtual switches group packets into flows by processing and matching
   particular packet or frame header information, or by matching packets
   based on the input ports. Thus, a flow can be thought of as a
   sequence of packets that have the same set of header field values or
   have arrived on the same physical or logical port. Performance
   results can vary based on the parameters the vSwitch uses to match
   for a flow. The recommended flow classification parameters for any
   vSwitch performance tests are: the input port (physical or logical),
   the source MAC address, the destination MAC address, the source IP
   address, the destination IP address, and the Ethernet protocol type
   field (although classification may take place on other fields, such
   as source and destination transport port numbers). It is essential
   to increase the flow timeout time on a vSwitch before conducting any
   performance tests that do not intend to measure the flow setup time
   (see Section 3 of [RFC2889]). Normally, the first packet of a
   particular stream will install the flow in the virtual switch, which
   introduces additional latency; subsequent packets of the same flow
   are not subject to this latency if the flow is already installed on
   the vSwitch.

3.5. Benchmarks Using Baselines with Resource Isolation

   This outline describes the measurement of baselines with isolated
   resources at a high level, which is the intended approach at this
   time.

   1. Baselines:

       * Optional: Benchmark platform forwarding capability without a
          vSwitch or VNF for at least 72 hours (serves as a means of
          platform validation and a means to obtain the base performance
          for the platform in terms of its maximum forwarding rate and
          latency).

```
                                                             __
        +----------------------------------------------+   |  |
        |   +--------------------------------------+   |   |  |
        |   |                                      |   |   |  |
        |   |         Simple Forwarding App        |   |   | Host
        |   |                                      |   |   |  |
        |   +--------------------------------------+   |   |  |
        |   |                  NIC                 |   |   |  |
        +---+--------------------------------------+---+   __|
            ^                                    :
            |                                    |
            :                                    v
        +----------------------------------------------+
        |                                              |
        |              Traffic Generator               |
        |                                              |
        +----------------------------------------------+
```
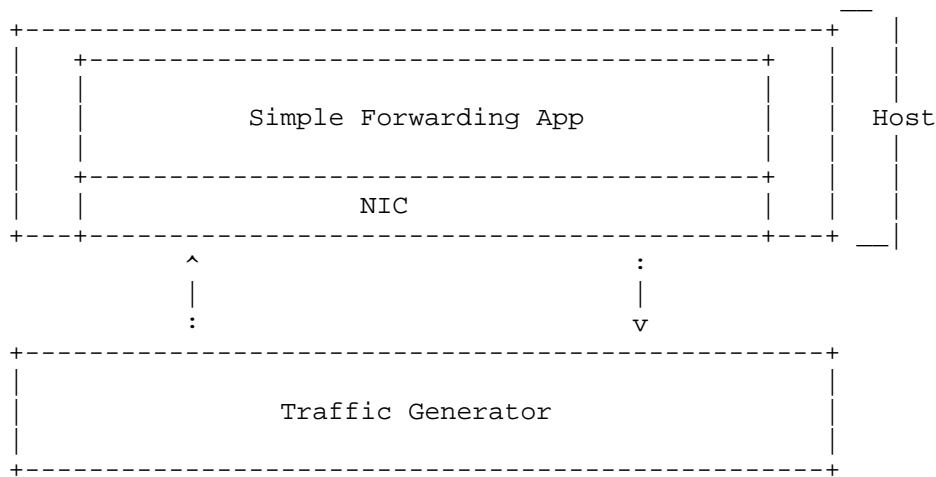
             Figure 1: Benchmark Platform Forwarding Capability

* Benchmark VNF forwarding capability with direct connectivity
  (vSwitch bypass, e.g., SR/IOV) for at least 72 hours (serves
  as a means of VNF validation and a means to obtain the base
  performance for the VNF in terms of its maximum forwarding
  rate and latency). The metrics gathered from this test will
  serve as a key comparison point for vSwitch bypass
  technologies performance and vSwitch performance.

```
                                                         __
  +-----------------------------------------------+ __  |  |
  |  +------------------------------------------+  |  | |  |
  |  |                                          |  |  | | Host/ |
  |  |                  VNF                     |  |  | | Guest |
  |  |                                          |  |  | |  |
  |  +------------------------------------------+  |  | |__| |
  |  |           Passthrough/SR-IOV             |  |  |     Host |
  |  +------------------------------------------+  |  |      |
  |  |                  NIC                     |  |  |      |
  +---+------------------------------------------+---+      __|
          ^                              :
          |                              |
          :                              v
  +-----------------------------------------------+
  |                                               |
  |              Traffic Generator                |
  |                                               |
  +-----------------------------------------------+
```
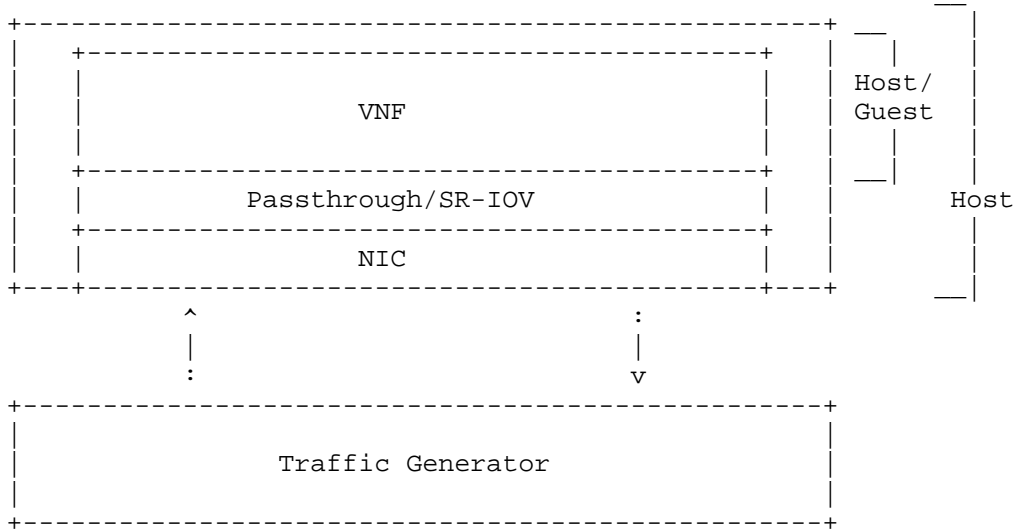
       Figure 2: Benchmark VNF Forwarding Capability

* Benchmarking with isolated resources alone and with other
  resources (both hardware and software) disabled; for example,
  vSwitch and VM are SUT.

* Benchmarking with isolated resources alone, thus leaving some
  resources unused.

* Benchmarking with isolated resources and all resources
  occupied.

2. Next Steps:

* Limited sharing

* Production scenarios

* Stressful scenarios

4. VSPERF Specification Summary

   The overall specification in preparation is referred to as a Level
   Test Design (LTD) document, which will contain a suite of performance
   tests. The base performance tests in the LTD are based on the
   pre-existing specifications developed by the BMWG to test the
   performance of physical switches. These specifications include:

   o Benchmarking Methodology for Network Interconnect Devices
      [RFC2544]

   o Benchmarking Methodology for LAN Switching [RFC2889]

   o Device Reset Characterization [RFC6201]

   o Packet Delay Variation Applicability Statement [RFC5481]

   The two most recent RFCs above ([RFC6201] and [RFC5481]) are being
   applied in benchmarking for the first time and represent a
   development challenge for test equipment developers. Fortunately,
   many members of the testing system community have engaged on the
   VSPERF project, including an open source test system.

   In addition to this, the LTD also reuses the terminology defined by:

   o Benchmarking Terminology for LAN Switching Devices [RFC2285]

   It is recommended that these references be included in future
   benchmarking specifications:

   o Methodology for IP Multicast Benchmarking [RFC3918]

   o Packet Reordering Metrics [RFC4737]

   As one might expect, the most fundamental internetworking
   characteristics of throughput and latency remain important when the
   switch is virtualized, and these benchmarks figure prominently in the
   specification.

   When considering characteristics important to "telco" network
   functions, additional performance metrics are needed. In this case,
   the project specifications have referenced metrics from the IETF IP
   Performance Metrics (IPPM) literature. This means that the latency
   test described in [RFC2544] is replaced by measurement of a metric
   derived from IPPM's [RFC7679], where a set of statistical summaries
   will be provided (mean, max, min, and percentiles). Further metrics
   planned to be benchmarked include packet delay variation as defined
   by [RFC5481], reordering, burst behaviour, DUT availability, DUT

capacity, and packet loss in long-term testing at the throughput
level, where some low level of background loss may be present and
characterized.

Tests have been designed to collect the metrics below:

o Throughput tests are designed to measure the maximum forwarding
   rate (in frames per second, fps) and bit rate (in Mbps) for a
   constant load (as defined by [RFC1242]) without traffic loss.

o Packet and frame-delay distribution tests are designed to measure
   the average minimum and maximum packet (and/or frame) delay for
   constant loads.

o Packet delay tests are designed to understand latency distribution
   for different packet sizes and to uncover outliers over an
   extended test run.

o Scalability tests are designed to understand how the virtual
   switch performs with an increasing number of flows, number of
   active ports, configuration complexity of the forwarding logic,
   etc.

o Stream performance tests (with TCP or UDP) are designed to measure
   bulk data transfer performance, i.e., how fast systems can send
   and receive data through the switch.

o Control-path and data-path coupling tests are designed to
   understand how closely the data path and the control path are
   coupled, as well as the effect of this coupling on the performance
   of the DUT (for example, delay of the initial packet of a flow).

o CPU and memory consumption tests are designed to understand the
   virtual switch's footprint on the system and are conducted as
   auxiliary measurements with the benchmarks above. They include
   CPU utilization, cache utilization, and memory footprint.

o The so-called "soak" tests, where the selected test is conducted
   over a long period of time (with an ideal duration of 24 hours but
   only long enough to determine that stability issues exist when
   found; there is no requirement to continue a test when a DUT
   exhibits instability over time). The key performance
   characteristics and benchmarks for a DUT are determined (using
   short duration tests) prior to conducting soak tests. The purpose
   of soak tests is to capture transient changes in performance,
   which may occur due to infrequent processes, memory leaks, or the
   low-probability coincidence of two or more processes. The
   stability of the DUT is the paramount consideration, so

performance must be evaluated periodically during continuous
testing, and this results in use of frame rate metrics [RFC2889]
instead of throughput [RFC2544] (which requires stopping traffic
to allow time for all traffic to exit internal queues), for
example.

Additional test specification development should include:

o Request/response performance tests (with TCP or UDP), which
   measure the transaction rate through the switch.

o Noisy neighbor tests, in order to understand the effects of
   resource sharing on the performance of a virtual switch.

o Tests derived from examination of ETSI NFV Draft GS IFA003
   requirements [IFA003] on characterization of acceleration
   technologies applied to vSwitches.

The flexibility of deployment of a virtual switch within a network
means that it is necessary to characterize the performance of a
vSwitch in various deployment scenarios. The deployment scenarios
under consideration are shown in the following figures:

```
                                                       __
     +-------------------------------------------------+   |
     |              +-------------------+              |   |
     |              |                   |              |   |
     |              |                   v              |   | Host
     |   +--------------+       +--------------+       |   |
     |   |   PHY Port   | vSwitch | PHY Port   |       |   |
     +---+--------------+----------+--------------+---+  __|
             ^                            :
             |                            |
             :                            v
     +-------------------------------------------------+
     |                                                 |
     |               Traffic Generator                 |
     |                                                 |
     +-------------------------------------------------+
```
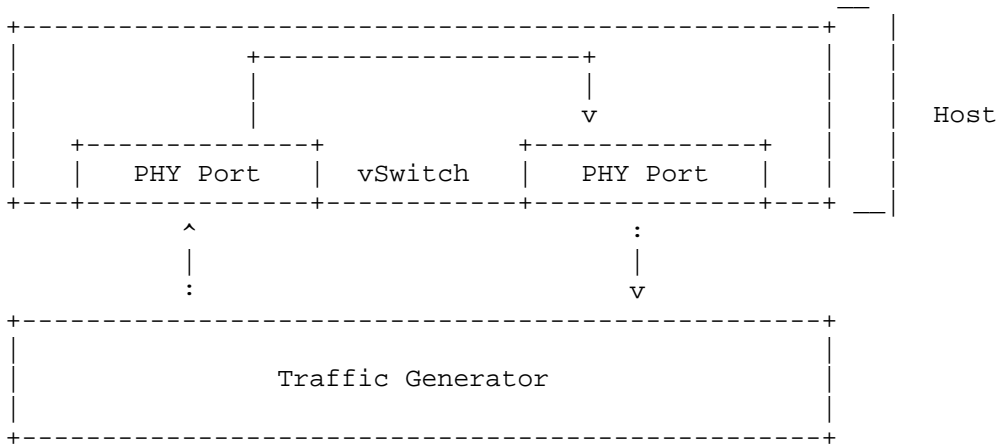
      Figure 3: Physical Port to Virtual Switch to Physical Port

```
                                                          __
       +-----------------------------------------------+  |
       |                                               |  |
       |   +---------------------------------------+   |  |
       |   |             Application               |   |  |
       |   +---------------------------------------+   |  |
       |       ^                           :           |  |
       |       |                           |           |  | Guest
       |       :                           v           |  |
       |   +--------------+       +--------------+      |  |
       |   | Logical Port 0|      | Logical Port 1|     |  |
       +---+--------------+-------+--------------+---+  __|
               ^                           :
               |                           |
               :                           v
       +---+--------------+-------+--------------+---+  __
       |   | Logical Port 0|      | Logical Port 1|   |  |
       |   +--------------+       +--------------+     |  |
       |       ^                           :           |  |
       |       |                           |           |  | Host
       |       :                           v           |  |
       |   +--------------+       +--------------+      |  |
       |   |  PHY Port    | vSwitch |  PHY Port  |      |  |
       +---+--------------+-------+--------------+---+  __|
               ^                           :
               |                           |
               :                           v
       +-----------------------------------------------+
       |                                               |
       |              Traffic Generator                |
       |                                               |
       +-----------------------------------------------+
```
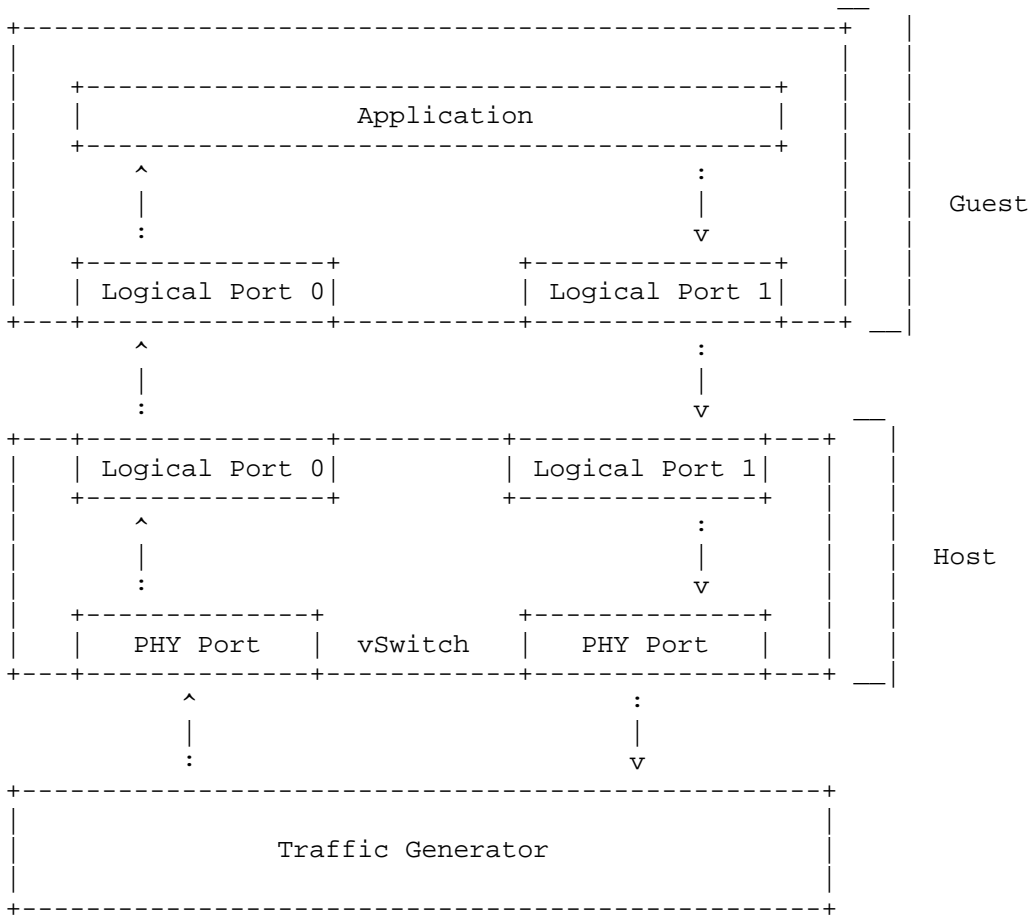
       Figure 4: Physical Port to Virtual Switch to VNF to Virtual Switch to
                             Physical Port

```
                                                          __
+--------------------+  +--------------------+  |
|   Guest 1          |  |   Guest 2          |  |
|   +--------------+ |  |   +--------------+ |  |
|   |  Application | |  |   |  Application | |  |
|   +--------------+ |  |   +--------------+ |  |
|         ^        | |  |         ^        | |  |
|         |        v |  |         |        v |  |      Guests
|   +--------------+ |  |   +--------------+ |  |
|   | Logical Ports| |  |   | Logical Ports| |  |
|   |  0        1  | |  |   |  0        1  | |  |
+---+--------------+--+  +---+--------------+--+__|
      ^          :            ^          :
      |          |            |          |
      :          v            :          v          _
+---+--------------+---------+--------------+--+  |
|   |  0        1  |         |  3        4  | |  |
|   | Logical Ports|         | Logical Ports| |  |
|   +--------------+         +--------------+ |  |
|         ^        |            ^          |  |  |      Host
|         |        \---------------/        v |  |
|   +--------------+         +--------------+ |  |
|   |  PHY Ports   | vSwitch |  PHY Ports   | |  |
+---+--------------+---------+--------------+---+_|
      ^                             :
      |                             |
      :                             v
+----------------------------------------------------+
|                                                    |
|               Traffic Generator                    |
|                                                    |
+----------------------------------------------------+
```
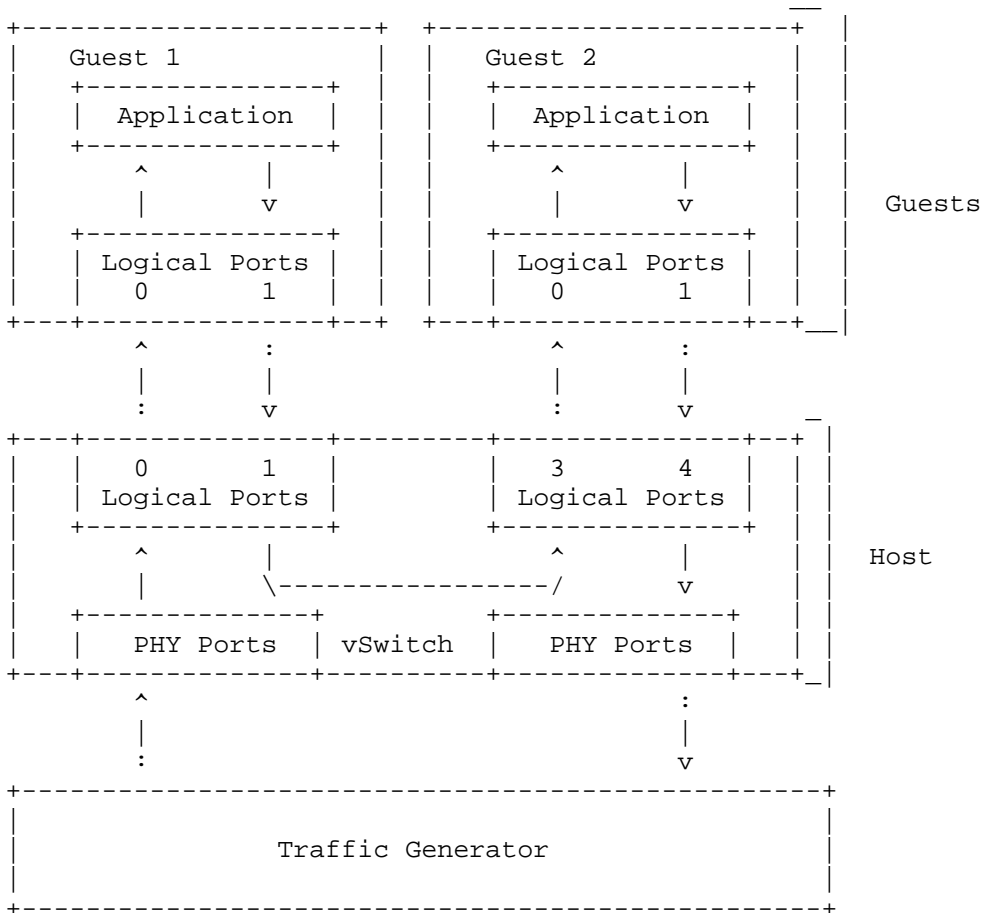
      Figure 5: Physical Port to Virtual Switch to VNF to Virtual Switch to
                VNF to Virtual Switch to Physical Port

```
                                                           __
   +----------------------------------------------------+   |
   |                                                    |   |
   |    +--------------------------------------------+  |   |
   |    |                 Application                 |  |   |
   |    +--------------------------------------------+  |   |
   |         ^                                          |   |
   |         |                                          |   |  Guest
   |         :                                          |   |
   |    +--------------+                                |   |
   |    | Logical Port 0|                               |   |
   +---+--------------+----------------------------+    |   __|
         ^                                              |
         |
         :
   +---+--------------+----------------------------+    |   __
   |    | Logical Port 0|                               |   |
   |    +--------------+                                |   |
   |         ^                                          |   |
   |         |                                          |   |  Host
   |         :                                          |   |
   |    +-------------+                                 |   |
   |    |   PHY Port  |    vSwitch                      |   |
   +---+-------------+----------- ------------- ---+    |   __|
         ^
         |
         :
   +----------------------------------------------------+
   |                                                    |
   |                 Traffic Generator                  |
   |                                                    |
   +----------------------------------------------------+
```
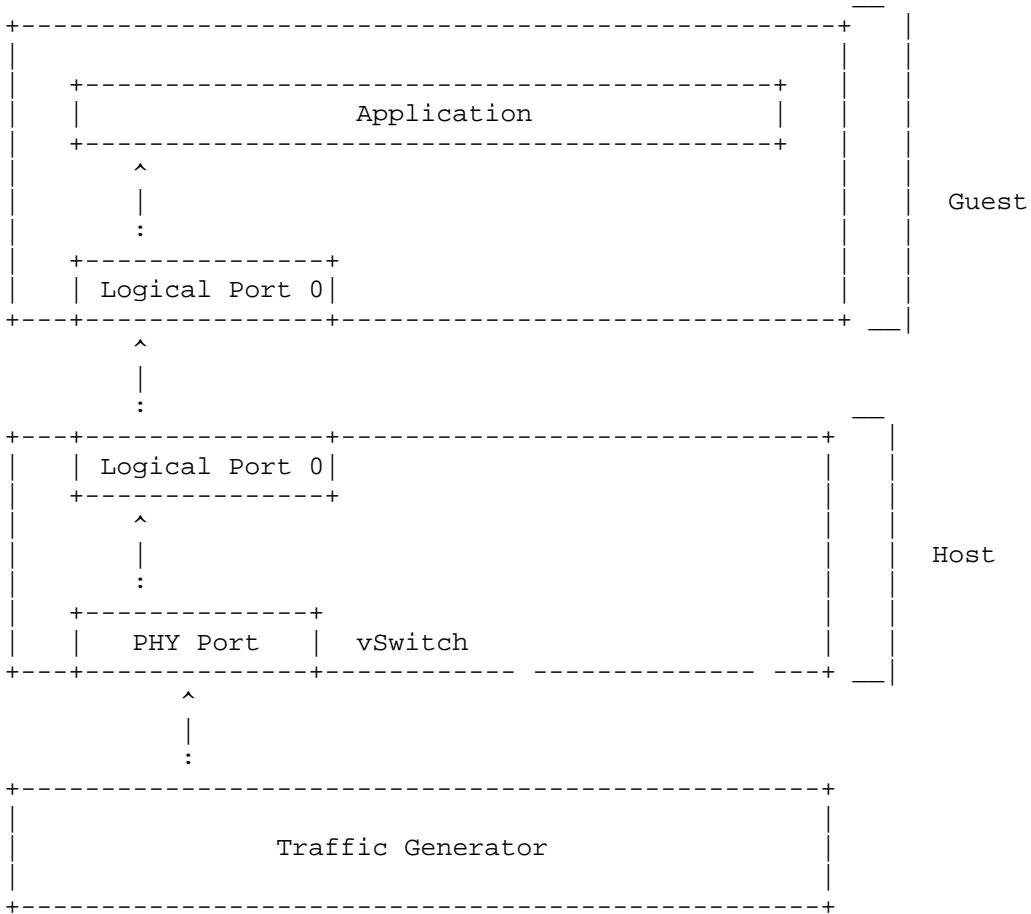
Figure 6: Physical Port to Virtual Switch to VNF

```
                                                     __
   +-----------------------------------------------+    |
   |                                               |    |  |
   |   +---------------------------------------+   |    |  |
   |   |              Application              |   |    |  |
   |   +---------------------------------------+   |    |  |
   |                                     :         |    |  |
   |                                     |         |    |  | Guest
   |                                     v         |    |  |
   |                            +--------------+   |    |  |
   |                            | Logical Port |   |    |  |
   +----------------------------+--------------+---+    __|
                                        :
                                        |
                                        v
                                                     __
   +----------------------------+--------------+---+    |
   |                            | Logical Port |   |    |  |
   |                            +--------------+   |    |  |
   |                                     :         |    |  |
   |                                     |         |    |  | Host
   |                                     v         |    |  |
   |                            +--------------+   |    |  |
   |    vSwitch                 |   PHY Port   |   |    |  |
   +----------------------------+--------------+---+    __|
                                        :
                                        |
                                        v
   +-----------------------------------------------+
   |                                               |
   |              Traffic Generator                |
   |                                               |
   +-----------------------------------------------+
```
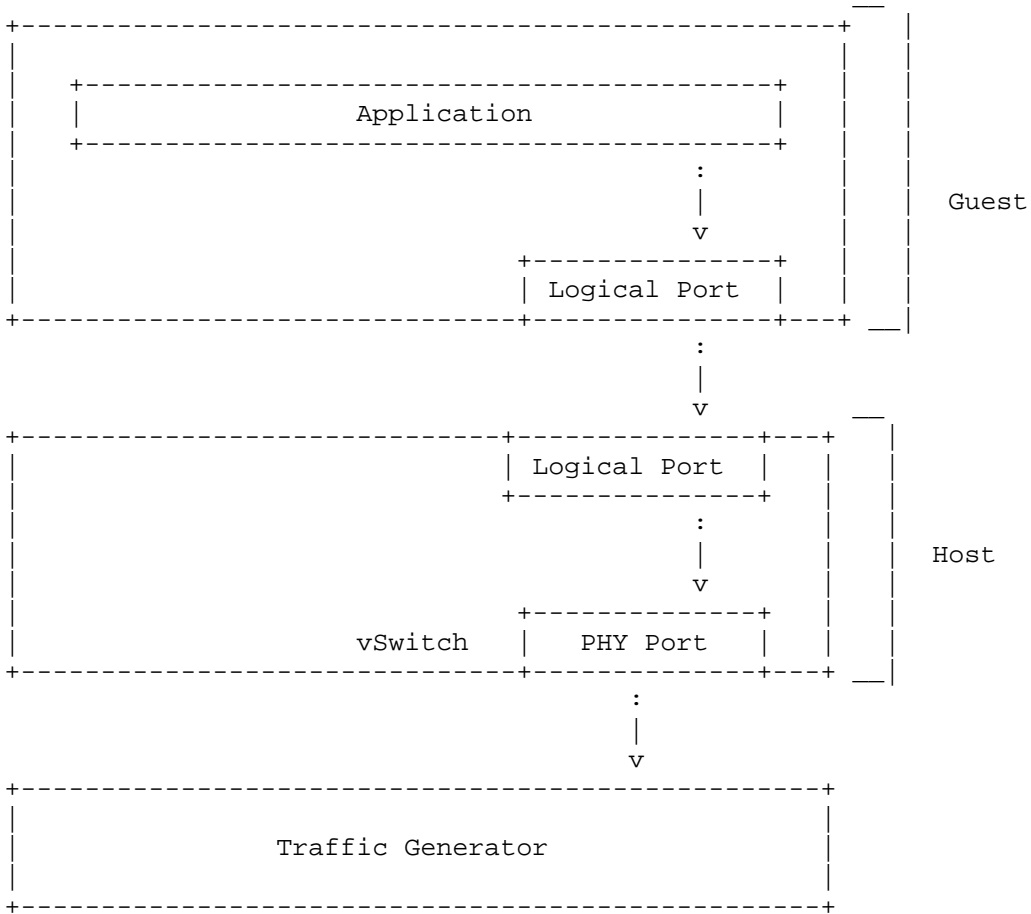
Figure 7: VNF to Virtual Switch to Physical Port

```
                                              __
   +--------------------+   +--------------------+  |
   |    Guest 1         |   |    Guest 2         |  |
   |   +--------------+ |   |   +--------------+ |  |
   |   |  Application | |   |   |  Application | |  |
   |   +--------------+ |   |   +--------------+ |  |
   |          |         |   |          ^         |  |
   |          v         |   |          |         |  |   Guests
   |   +--------------+ |   |   +--------------+ |  |
   |   | Logical Ports| |   |   | Logical Ports| |  |
   |   |            0  | |   |   |            0 | |  |
   +---+--------------+-+   +---+--------------+-+__|
           :                          ^
           |                          |
           v                          :            _
   +---+--------------+--------+--------------+--+  |
   |   |            1 |        |            1 |  |  |
   |   | Logical Ports|        | Logical Ports|  |  |
   |   +--------------+        +--------------+  |  |
   |          |                       ^          |  |   Host
   |          \---------------/                  |  |
   |                                             |  |
   |                 vSwitch                     |  |
   +---------------------------------------------+__|
```
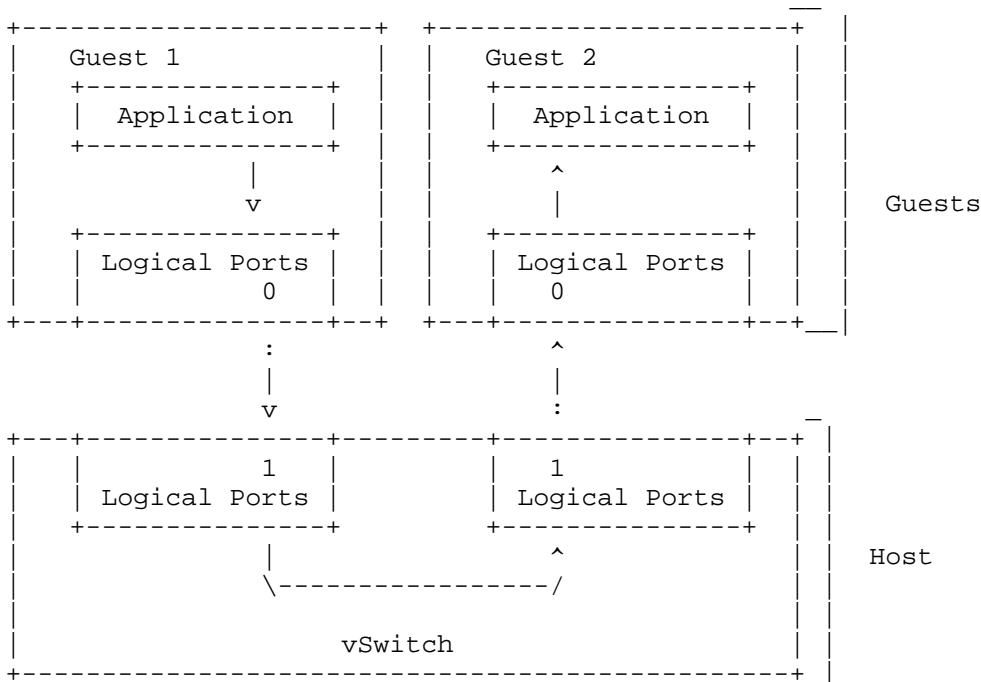
                 Figure 8: VNF to Virtual Switch to VNF

   A set of deployment scenario figures is available on the VSPERF "Test
   Methodology" wiki page [TestTopo].

5. 3x3 Matrix Coverage

   This section organizes the many existing test specifications into the
   "3x3" matrix (introduced in [RFC8172]). Because the LTD
   specification ID names are quite long, this section is organized into
   lists for each occupied cell of the matrix (not all are occupied;
   also, the matrix has grown to 3x4 to accommodate scale metrics when
   displaying the coverage of many metrics/benchmarks). The current
   version of the LTD specification is available; see [LTD].

   The tests listed below assess the activation of paths in the data
   plane rather than the control plane.

   A complete list of tests with short summaries is available on the
   VSPERF "LTD Test Spec Overview" wiki page [LTDoverV].

5.1. Speed of Activation

   o Activation.RFC2889.AddressLearningRate

   o PacketLatency.InitialPacketProcessingLatency

5.2. Accuracy of Activation

   o CPDP.Coupling.Flow.Addition

5.3. Reliability of Activation

   o Throughput.RFC2544.SystemRecoveryTime

   o Throughput.RFC2544.ResetTime

5.4. Scale of Activation

   o Activation.RFC2889.AddressCachingCapacity

5.5. Speed of Operation

   o Throughput.RFC2544.PacketLossRate

   o Stress.RFC2544.0PacketLoss

   o Throughput.RFC2544.PacketLossRateFrameModification

   o Throughput.RFC2544.BackToBackFrames

   o Throughput.RFC2889.MaxForwardingRate

   o Throughput.RFC2889.ForwardPressure

   o Throughput.RFC2889.BroadcastFrameForwarding

   o Throughput.RFC2544.WorstN-BestN

   o Throughput.Overlay.Network.<tech>.RFC2544.PacketLossRatio

5.6. Accuracy of Operation

   o Throughput.RFC2889.ErrorFramesFiltering

   o Throughput.RFC2544.Profile

5.7. Reliability of Operation

   o Throughput.RFC2889.Soak

   o Throughput.RFC2889.SoakFrameModification

   o PacketDelayVariation.RFC3393.Soak

5.8. Scalability of Operation

   o Scalability.RFC2544.0PacketLoss

   o MemoryBandwidth.RFC2544.0PacketLoss.Scalability

   o Scalability.VNF.RFC2544.PacketLossProfile

   o Scalability.VNF.RFC2544.PacketLossRatio

5.9. Summary

```
   |---------------------------------------------------------------|
   |              |         |          |             |             |
   |              |  SPEED  | ACCURACY | RELIABILITY |    SCALE    |
   |              |         |          |             |             |
   |---------------------------------------------------------------|
   |              |         |          |             |             |
   |  Activation  |    X    |    X     |      X      |      X      |
   |              |         |          |             |             |
   |---------------------------------------------------------------|
   |              |         |          |             |             |
   |  Operation   |    X    |    X     |      X      |      X      |
   |              |         |          |             |             |
   |---------------------------------------------------------------|
   |              |         |          |             |             |
   | De-activation|         |          |             |             |
   |              |         |          |             |             |
   |---------------------------------------------------------------|
```

6. Security Considerations

   Benchmarking activities as described in this memo are limited to
   technology characterization of a Device Under Test/System Under Test
   (DUT/SUT) using controlled stimuli in a laboratory environment with
   dedicated address space and the constraints specified in the sections
   above.

   The benchmarking network topology will be an independent test setup
   and MUST NOT be connected to devices that may forward the test
   traffic into a production network or misroute traffic to the test
   management network.

   Further, benchmarking is performed on a "black-box" basis and relies
   solely on measurements observable external to the DUT/SUT.

   Special capabilities SHOULD NOT exist in the DUT/SUT specifically for
   benchmarking purposes. Any implications for network security arising
   from the DUT/SUT SHOULD be identical in the lab and in production
   networks.

7. References

7.1. Normative References

   [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
              Requirement Levels", BCP 14, RFC 2119,
              DOI 10.17487/RFC2119, March 1997,
              <https://www.rfc-editor.org/info/rfc2119>.

   [RFC2285] Mandeville, R., "Benchmarking Terminology for LAN
              Switching Devices", RFC 2285, DOI 10.17487/RFC2285,
              February 1998, <https://www.rfc-editor.org/info/rfc2285>.

   [RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for
              Network Interconnect Devices", RFC 2544,
              DOI 10.17487/RFC2544, March 1999,
              <https://www.rfc-editor.org/info/rfc2544>.

   [RFC2889] Mandeville, R. and J. Perser, "Benchmarking Methodology
              for LAN Switching Devices", RFC 2889,
              DOI 10.17487/RFC2889, August 2000,
              <https://www.rfc-editor.org/info/rfc2889>.

   [RFC3918] Stopp, D. and B. Hickman, "Methodology for IP Multicast
              Benchmarking", RFC 3918, DOI 10.17487/RFC3918, October
              2004, <https://www.rfc-editor.org/info/rfc3918>.

   [RFC4737] Morton, A., Ciavattone, L., Ramachandran, G., Shalunov,
              S., and J. Perser, "Packet Reordering Metrics", RFC 4737,
              DOI 10.17487/RFC4737, November 2006,
              <https://www.rfc-editor.org/info/rfc4737>.

   [RFC6201] Asati, R., Pignataro, C., Calabria, F., and C. Olvera,
              "Device Reset Characterization", RFC 6201,
              DOI 10.17487/RFC6201, March 2011,
              <https://www.rfc-editor.org/info/rfc6201>.

   [RFC6985] Morton, A., "IMIX Genome: Specification of Variable Packet
              Sizes for Additional Testing", RFC 6985,
              DOI 10.17487/RFC6985, July 2013,
              <https://www.rfc-editor.org/info/rfc6985>.

   [RFC7679] Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton,
              Ed., "A One-Way Delay Metric for IP Performance Metrics
              (IPPM)", STD 81, RFC 7679, DOI 10.17487/RFC7679, January
              2016, <https://www.rfc-editor.org/info/rfc7679>.

   [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC
              2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174,
              May 2017, <https://www.rfc-editor.org/info/rfc8174>.

7.2. Informative References

   [BENCHMARK-METHOD]
              Huang, L., Ed., Rong, G., Ed., Mandeville, B., and B.
              Hickman, "Benchmarking Methodology for Virtualization
              Network Performance", Work in Progress, draft-huang-bmwg-
              virtual-network-performance-03, July 2017.

   [DanubeRel]
              OPNFV, "Danube",
              <https://wiki.opnfv.org/display/SWREL/Danube>.

   [IEEE802.1ac]
              IEEE, "IEEE Standard for Local and metropolitan area
              networks -- Media Access Control (MAC) Service
              Definition", IEEE 802.1AC-2016,
              DOI 10.1109/IEEESTD.2017.7875381, 2016,
              <https://standards.ieee.org/findstds/
              standard/802.1AC-2016.html>.

   [IEEE829] IEEE, "IEEE Standard for Software and System Test
              Documentation", IEEE 829-2008,
              DOI 10.1109/IEEESTD.2008.4578383,
              <http://ieeexplore.ieee.org/document/4578383/>.

   [IFA003]   ETSI, "Network Functions Virtualisation (NFV);
              Acceleration Technologies; vSwitch Benchmarking and
              Acceleration Specification", ETSI GS NFV-IFA 003 V2.1.1,
              April 2016, <http://www.etsi.org/deliver/etsi_gs/
              NFV-IFA/001_099/003/02.01.01_60/
              gs_NFV-IFA003v020101p.pdf>.

   [LTD]   Tahhan, M., "VSPERF Level Test Design (LTD)",
              <http://docs.opnfv.org/en/stable-danube/
              submodules/vswitchperf/docs/testing/developer/
              requirements/vswitchperf_ltd.html#>.

   [LTDoverV] Morton, A., "LTD Test Spec Overview",
              <https://wiki.opnfv.org/display/vsperf/
              LTD+Test+Spec+Overview>.

   [OPNFV]   OPNFV, "OPNFV", <https://www.opnfv.org/>.

   [RFC1242] Bradner, S., "Benchmarking Terminology for Network
              Interconnection Devices", RFC 1242, DOI 10.17487/RFC1242,
              July 1991, <https://www.rfc-editor.org/info/rfc1242>.

   [RFC5481] Morton, A. and B. Claise, "Packet Delay Variation
              Applicability Statement", RFC 5481, DOI 10.17487/RFC5481,
              March 2009, <https://www.rfc-editor.org/info/rfc5481>.

   [RFC8172] Morton, A., "Considerations for Benchmarking Virtual
              Network Functions and Their Infrastructure", RFC 8172,
              DOI 10.17487/RFC8172, July 2017,
              <https://www.rfc-editor.org/info/rfc8172>.

   [TestTopo] Snyder, E., "Test Methodology",
              <https://wiki.opnfv.org/display/vsperf/Test+Methodology>.

   [VSPERFhome]
              Tahhan, M., "VSPERF Home",
              <https://wiki.opnfv.org/display/vsperf/VSperf+Home>.

Authors' Addresses

    Maryam Tahhan
    Intel

    Email: maryam.tahhan@intel.com

    Billy O'Mahony
    Intel

    Email: billy.o.mahony@intel.com

    Al Morton
    AT&T Labs
    200 Laurel Avenue South
    Middletown, NJ 07748
    United States of America

    Phone: +1 732 420 1571
    Fax: +1 732 368 1192
    Email: acmorton@att.com