



Yardstick Overview

Release draft (93fb39a)

OPNFV

January 08, 2016

CONTENTS

1	Introduction	1
1.1	Contact Yardstick	1
2	Methodology	3
2.1	Abstract	3
2.2	ETSI-NFV	3
2.3	Metrics	3
3	Yardstick Test Cases	7
3.1	Abstract	7
3.2	Generic NFVI Test Case Descriptions	8
3.3	OPNFV Feature Test Cases	14
3.4	Templates	14
4	Yardstick Glossary	17
	Index	19

INTRODUCTION

Welcome to Yardstick's documentation !

[Yardstick](#) is an OPNFV Project.

The project's goal is to verify infrastructure compliance, from the perspective of a *VNF*.

The Project's scope is the development of a test framework, *Yardstick*, test cases and test stimuli to enable *NFVI* verification. The Project also includes a sample *VNF*, the *VTC* and its experimental framework, *ApexLake* !

The chapter [Methodology](#) describes the methodology implemented by the Yardstick Project for *NFVI* verification. The chapter [Yardstick Test Cases](#) includes a list of available Yardstick test cases.

Yardstick is used for verifying the OPNFV infrastructure and some of the OPNFV features, listed in [Yardstick Test Cases](#).

The *Yardstick* framework is deployed in several OPNFV community labs. It is installer, infrastructure and application independent.

See also:

[Pharos](#) for information on OPNFV community labs.

1.1 Contact Yardstick

Feedback? [Contact us](#)

METHODOLOGY

2.1 Abstract

This chapter describes the methodology implemented by the Yardstick project for verifying the NFV Infrastructure from the perspective of a VNF.

2.2 ETSI-NFV

The document ETSI GS [NFV-TST001](#), “Pre-deployment Testing; Report on Validation of NFV Environments and Services”, recommends methods for pre-deployment testing of the functional components of an NFV environment.

The Yardstick project implements the methodology described in chapter 6, “Pre- deployment validation of NFV infrastructure”.

The methodology consists in decomposing the typical VNF work-load performance metrics into a number of characteristics/performance vectors, which each can be represented by distinct test-cases.

The methodology includes five steps:

- **Step1: Define Infrastructure - the HW, SW and corresponding configuration** target for validation; the OPNFV infrastructure, in OPNFV community labs.
- **Step2: Identify VNF type - the application for which the infrastructure is** to be validated, and its requirements on the underlying infrastructure.
- **Step3: Select test cases - depending on the workload that represents the** application for which the infrastructure is to be validated, the relevant test cases amongst the list of available Yardstick test cases.
- **Step4: Execute tests - define the duration and number of iterations for the** selected test cases, tests runs are automated via OPNFV Jenkins Jobs.
- **Step5:** Collect results - using the common API for result collection.

2.3 Metrics

The metrics, as defined by ETSI GS [NFV-TST001](#), are shown in [Table1](#), [Table2](#) and [Table3](#).

In OPNFV Brahmaputra release, generic test cases covering aspects of the listed metrics are available; further OPNFV releases will provide extended testing of these metrics. The view of available Yardstick test cases cross ETSI definitions in [Table1](#), [Table2](#) and [Table3](#) is shown in [Table4](#). It shall be noticed that the Yardstick test cases are examples, the test duration and number of iterations are configurable, as are the System Under Test (SUT) and the attributes (or, in Yardstick nomenclature, the scenario options). **Table 1 - Performance/Speed Metrics**

Category	Performance/Speed
Compute	<ul style="list-style-type: none"> • Latency for random memory access • Latency for cache read/write operations • Processing speed (instructions per second) • Throughput for random memory access (bytes per second)
Network	<ul style="list-style-type: none"> • Throughput per NFVI node (frames/byte per second) • Throughput provided to a VM (frames/byte per second) • Latency per traffic flow • Latency between VMs • Latency between NFVI nodes • Packet delay variation (jitter) between VMs • Packet delay variation (jitter) between NFVI nodes
Storage	<ul style="list-style-type: none"> • Sequential read/write IOPS • Random read/write IOPS • Latency for storage read/write operations • Throughput for storage read/write operations

Table 2 - Capacity/Scale Metrics

Category	Capacity/Scale
Compute	<ul style="list-style-type: none"> • Number of cores and threads- Available memory size • Cache size • Processor utilization (max, average, standard deviation) • Memory utilization (max, average, standard deviation) • Cache utilization (max, average, standard deviation)
Network	<ul style="list-style-type: none"> • Number of connections • Number of frames sent/received • Maximum throughput between VMs (frames/byte per second) • Maximum throughput between NFVI nodes (frames/byte per second) • Network utilization (max, average, standard deviation) • Number of traffic flows
Storage	<ul style="list-style-type: none"> • Storage/Disk size • Capacity allocation (block-based, object-based) • Block size • Maximum sequential read/write IOPS • Maximum random read/write IOPS • Disk utilization (max, average, standard deviation)

Table 3 - Availability/Reliability Metrics

Category	Availability/Reliability
Compute	<ul style="list-style-type: none"> • Processor availability (Error free processing time) • Memory availability (Error free memory time) • Processor mean-time-to-failure • Memory mean-time-to-failure • Number of processing faults per second
Network	<ul style="list-style-type: none"> • NIC availability (Error free connection time) • Link availability (Error free transmission time) • NIC mean-time-to-failure • Network timeout duration due to link failure • Frame loss rate
Storage	<ul style="list-style-type: none"> • Disk availability (Error free disk access time) • Disk mean-time-to-failure • Number of failed storage read/write operations per second

Table 4 - Yardstick Generic Test Cases

Category	Performance/Speed	Capacity/Scale	Availability/Reliability
Compute	TC003 TC004 TC014 TC024	TC003 TC004 TC010 TC012	TC013 ¹ TC015 ¹
Network	TC002 TC011	TC001 TC008 TC009	TC016 ¹ TC018 ¹
Storage	TC005	TC005	TC017 ¹

Note: The description in this OPNFV document is intended as a reference for users to understand the scope of the Yardstick Project and the deliverables of the Yardstick framework. For complete description of the methodology, refer to the ETSI document.

¹To be included in future deliveries.

YARDSTICK TEST CASES

3.1 Abstract

This chapter lists available Yardstick test cases. Yardstick test cases are divided in two main categories:

- *Generic NFVI Test Cases* - Test Cases developed to realize the methodology described in [Methodology](#)
- *OPNFV Feature Test Cases* - Test Cases developed to verify one or more aspect of a feature delivered by an OPNFV Project.

3.2 Generic NFVI Test Case Descriptions

3.2.1 Yardstick Test Case Description TC001

Network Performance	
test case id	OPNFV_YARDSTICK_TC001_NW PERF
metric	Number of flows and throughput
test purpose	To evaluate the IaaS network performance with regards to flows and throughput, such as if and how different amounts of flows matter for the throughput between hosts on different compute blades. Typically e.g. the performance of a vSwitch depends on the number of flows running through it. Also performance of other equipment or entities can depend on the number of flows or the packet sizes used. The purpose is also to be able to spot trends. Test results, graphs and similar shall be stored for comparison reasons and product evolution understanding between different OPNFV versions and/or configurations.
configuration	file: opnfv_yardstick_tc001.yaml Packet size: 60 bytes Number of ports: 10, 50, 100, 500 and 1000, where each runs for 20 seconds. The whole sequence is run twice. The client and server are distributed on different HW. For SLA max_ppm is set to 1000. The amount of configured ports map to between 110 up to 1001000 flows, respectively.
test tool	pktgen (Pktgen is not always part of a Linux distribution, hence it needs to be installed. It is part of the Yardstick Docker image. As an example see the /yardstick/tools/ directory for how to generate a Linux image with pktgen included.)
references	pktgen ETSI-NFV-TST001
applicability	Test can be configured with different packet sizes, amount of flows and test duration. Default values exist. SLA (optional): max_ppm: The number of packets per million packets sent that are acceptable to loose, not received.
pre-test conditions	The test case image needs to be installed into Glance with pktgen included in it. No POD specific requirements have been identified.
test sequence	description and expected result
step 1	The hosts are installed, as server and client. pktgen is invoked and logs are produced and stored. Result: Logs are stored.
test verdict	Fails only if SLA is not passed, or if there is a test case execution problem.

3.2.2 Yardstick Test Case Description TC002

Network Latency	
test case id	OPNFV_YARDSTICK_TC002_NW_LATENCY
metric	RTT, Round Trip Time
test purpose	To do a basic verification that network latency is within acceptable boundaries when packets travel between hosts located on same or different compute blades. The purpose is also to be able to spot trends. Test results, graphs and similar shall be stored for comparison reasons and product evolution understanding between different OPNFV versions and/or configurations.
configuration	file: opnfv_yardstick_tc002.yaml Packet size 100 bytes. Total test duration 600 seconds. One ping each 10 seconds. SLA RTT is set to maximum 10 ms.
test tool	ping Ping is normally part of any Linux distribution, hence it doesn't need to be installed. It is also part of the Yardstick Docker image. (For example also a Cirros image can be downloaded from cirros-image , it includes ping)
references	Ping man page ETSI-NFV-TST001
applicability	Test case can be configured with different packet sizes, burst sizes, ping intervals and test duration. SLA is optional. The SLA in this test case serves as an example. Considerably lower RTT is expected, and also normal to achieve in balanced L2 environments. However, to cover most configurations, both bare metal and fully virtualized ones, this value should be possible to achieve and acceptable for black box testing. Many real time applications start to suffer badly if the RTT time is higher than this. Some may suffer bad also close to this RTT, while others may not suffer at all. It is a compromise that may have to be tuned for different configuration purposes.
pre-test conditions	The test case image needs to be installed into Glance with ping included in it. No POD specific requirements have been identified.
test sequence	description and expected result
step 1	The hosts are installed, as server and client. Ping is invoked and logs are produced and stored. Result: Logs are stored.
test verdict	Test should not PASS if any RTT is above the optional SLA value, or if there is a test case execution problem.

3.2.3 Yardstick Test Case Description TC008

Packet Loss Extended Test	
test case id	OPNFV_YARDSTICK_TC008_NW PERF, Packet loss Extended Test
metric	Number of flows, packet size and throughput
test purpose	To evaluate the IaaS network performance with regards to flows and throughput, such as if and how different amounts of packet sizes and flows matter for the throughput between VMs on different compute blades. Typically e.g. the performance of a vSwitch depends on the number of flows running through it. Also performance of other equipment or entities can depend on the number of flows or the packet sizes used. The purpose is also to be able to spot trends. Test results, graphs and similar shall be stored for comparison reasons and product evolution understanding between different OPNFV versions and/or configurations.
configuration	file: opnfv_yardstick_tc008.yaml Packet size: 64, 128, 256, 512, 1024, 1280 and 1518 bytes. Number of ports: 1, 10, 50, 100, 500 and 1000. The amount of configured ports map from 2 up to 1001000 flows, respectively. Each packet_size/port_amount combination is run ten times, for 20 seconds each. Then the next packet_size/port_amount combination is run, and so on. The client and server are distributed on different HW. For SLA max_ppm is set to 1000.
test tool	pktgen (Pktgen is not always part of a Linux distribution, hence it needs to be installed. It is part of the Yardstick Docker image. As an example see the /yardstick/tools/ directory for how to generate a Linux image with pktgen included.)
references	pktgen ETSI-NFV-TST001
applicability	Test can be configured with different packet sizes, amount of flows and test duration. Default values exist. SLA (optional): max_ppm: The number of packets per million packets sent that are acceptable to loose, not received.
pre-test conditions	The test case image needs to be installed into Glance with pktgen included in it. No POD specific requirements have been identified.
test sequence	description and expected result
step 1	The hosts are installed, as server and client. pktgen is invoked and logs are produced and stored. Result: Logs are stored.
test verdict	Fails only if SLA is not passed, or if there is a test case execution problem.

3.2.4 Yardstick Test Case Description TC009

Packet Loss	
test case id	OPNFV_YARDSTICK_TC009_NW PERF, Packet loss
metric	Number of flows, packets lost and throughput
test purpose	To evaluate the IaaS network performance with regards to flows and throughput, such as if and how different amounts of flows matter for the throughput between VMs on different compute blades. Typically e.g. the performance of a vSwitch depends on the number of flows running through it. Also performance of other equipment or entities can depend on the number of flows or the packet sizes used. The purpose is also to be able to spot trends. Test results, graphs and similar shall be stored for comparison reasons and product evolution understanding between different OPNFV versions and/or configurations.
configuration	file: opnfv_yardstick_tc009.yaml Packet size: 64 bytes Number of ports: 1, 10, 50, 100, 500 and 1000. The amount of configured ports map from 2 up to 1001000 flows, respectively. Each port amount is run ten times, for 20 seconds each. Then the next port_amount is run, and so on. The client and server are distributed on different HW. For SLA max_ppm is set to 1000.
test tool	pktgen (Pktgen is not always part of a Linux distribution, hence it needs to be installed. It is part of the Yardstick Docker image. As an example see the /yardstick/tools/ directory for how to generate a Linux image with pktgen included.)
references	pktgen ETSI-NFV-TST001
applicability	Test can be configured with different packet sizes, amount of flows and test duration. Default values exist. SLA (optional): max_ppm: The number of packets per million packets sent that are acceptable to loose, not received.
pre-test conditions	The test case image needs to be installed into Glance with pktgen included in it. No POD specific requirements have been identified.
test sequence	description and expected result
step 1	The hosts are installed, as server and client. pktgen is invoked and logs are produced and stored. Result: logs are stored.
test verdict	Fails only if SLA is not passed, or if there is a test case execution problem.

3.2.5 Yardstick Test Case Description TC010

Memory Latency	
test case id	OPNFV_YARDSTICK_TC010_Memory Latency
metric	Latency in nanoseconds
test purpose	Measure the memory read latency for varying memory sizes and strides. Whole memory hierarchy is measured including all levels of cache.
configuration	File: opnfv_yardstick_tc010.yaml <ul style="list-style-type: none"> • SLA (max_latency): 30 nanoseconds • Stride - 128 bytes • Stop size - 64 megabytes • Iterations: 10 - test is run 10 times iteratively. • Interval: 1 - there is 1 second delay between each iteration.
test tool	Lmbench Lmbench is a suite of operating system microbenchmarks. This test uses lat_mem_rd tool from that suite. Lmbench is not always part of a Linux distribution, hence it needs to be installed in the test image
references	man-pages McVoy, Larry W.,and Carl Staelin. “Imbench: Portable Tools for Performance Analysis.” USENIX annual technical conference 1996.
applicability	Test can be configured with different: <ul style="list-style-type: none"> • strides; • stop_size; • iterations and intervals. There are default values for each above-mentioned option. SLA (optional) : max_latency: The maximum memory latency that is accepted.
pre-test conditions	The test case image needs to be installed into Glance with Lmbench included in the image. No POD specific requirements have been identified.
test sequence	description and expected result
step 1	The host is installed as client. Lmbench’s lat_mem_rd tool is invoked and logs are produced and stored. Result: logs are stored.
test verdict	Test fails if the measured memory latency is above the SLA value or if there is a test case execution problem.

3.2.6 Yardstick Test Case Description TC012

Memory Bandwidth	
test case id	OPNFV_YARDSTICK_TC012_Memory Bandwidth
metric	Megabyte per second (MBps)
test purpose	Measure the rate at which data can be read from and written to the memory (this includes all levels of memory).
configuration	<p>File: opnfv_yardstick_tc012.yaml</p> <ul style="list-style-type: none"> • SLA (optional): 15000 (MBps) min_bw: The minimum amount of memory bandwidth that is accepted. • Size: 10 240 kB - test allocates twice that size (20 480kB) zeros it and then measures the time it takes to copy from one side to another. • Benchmark: rdwr - measures the time to read data into memory and then write data to the same location. • Warmup: 0 - the number of iterations to perform before taking actual measurements. • Iterations: 10 - test is run 10 times iteratively. • Interval: 1 - there is 1 second delay between each iteration.
test tool	<p>Lmbench</p> <p>Lmbench is a suite of operating system microbenchmarks. This test uses bw_mem tool from that suite. Lmbench is not always part of a Linux distribution, hence it needs to be installed in the test image.</p>
references	<p>man-pages</p> <p>McVoy, Larry W., and Carl Staelin. "lmbench: Portable Tools for Performance Analysis." USENIX annual technical conference. 1996.</p>
applicability	<p>Test can be configured with different:</p> <ul style="list-style-type: none"> • memory sizes; • memory operations (such as rd, wr, rdwr, cp, frd, fwr, fcp, bzero, bcopy); • number of warmup iterations; • iterations and intervals. <p>There are default values for each above-mentioned option.</p>
pre-test conditions	<p>The test case image needs to be installed into Glance with Lmbench included in the image.</p> <p>No POD specific requirements have been identified.</p>
test sequence	description and expected result
step 1	<p>The host is installed as client. Lmbench's bw_mem tool is invoked and logs are produced and stored.</p> <p>Result: logs are stored.</p>
test verdict	<p>Test fails if the measured memory bandwidth is below the SLA value or if there is a test case execution problem.</p>

3.3 OPNFV Feature Test Cases

3.4 Templates

3.4.1 Yardstick Test Case Description TCXXX

test case slogan e.g. Network Latency	
test case id	e.g. OPNFV_YARDSTICK_TC001_NW Latency
metric	what will be measured, e.g. latency
test purpose	describe what is the purpose of the test case
configuration	what .yaml file to use, state SLA if applicable, state test duration, list and describe the scenario options used in this TC and also list the options using default values.
test tool	e.g. ping
references	e.g. RFCxxx, ETSI-NFVyyy
applicability	describe variations of the test case which can be performed, e.g. run the test for different packet sizes
pre-test conditions	describe configuration in the tool(s) used to perform the measurements (e.g. fio, pktgen), POD-specific configuration required to enable running the test
test sequence	description and expected result
step 1	use this to describe tests that require several steps e.g collect logs. Result: what happens in this step e.g. logs collected
step 2	remove interface Result: interface down.
step N	what is done in step N Result: what happens
test verdict	expected behavior, or SLA, pass/fail criteria

3.4.2 Task Template Syntax

Basic template syntax

A nice feature of the input task format used in Yardstick is that it supports the template syntax based on Jinja2. This turns out to be extremely useful when, say, you have a fixed structure of your task but you want to parameterize this task in some way. For example, imagine your input task file (task.yaml) runs a set of Ping scenarios:

```
# Sample benchmark task config file
# measure network latency using ping
schema: "yardstick:task:0.1"

scenarios:
-
  type: Ping
  options:
    packetsize: 200
    host: athena.demo
    target: ares.demo

  runner:
    type: Duration
    duration: 60
```

```

interval: 1

sla:
  max_rtt: 10
  action: monitor

context:
  ...

```

Let's say you want to run the same set of scenarios with the same runner/ context/sla, but you want to try another packetsize to compare the performance. The most elegant solution is then to turn the packetsize name into a template variable:

```

# Sample benchmark task config file
# measure network latency using ping

schema: "yardstick:task:0.1"
scenarios:
-
  type: Ping
  options:
    packetsize: {{packetsize}}
  host: athena.demo
  target: ares.demo

  runner:
    type: Duration
    duration: 60
    interval: 1

  sla:
    max_rtt: 10
    action: monitor

context:
  ...

```

and then pass the argument value for `{{packetsize}}` when starting a task with this configuration file. Yardstick provides you with different ways to do that:

1.Pass the argument values directly in the command-line interface (with either a JSON or YAML dictionary):

```

yardstick task start samples/ping-template.yaml
--task-args '{"packetsize": "200"}'

```

2.Refer to a file that specifies the argument values (JSON/YAML):

```

yardstick task start samples/ping-template.yaml --task-args-file args.yaml

```

Using the default values

Note that the Jinja2 template syntax allows you to set the default values for your parameters. With default values set, your task file will work even if you don't parameterize it explicitly while starting a task. The default values should be set using the `{% set ... %}` clause (task.yaml). For example:

```

# Sample benchmark task config file
# measure network latency using ping
schema: "yardstick:task:0.1"

```

```
{% set packetsize = packetsize or "100" %}
scenarios:
-
  type: Ping
  options:
  packetsize: {{packetsize}}
  host: athena.demo
  target: ares.demo

  runner:
    type: Duration
    duration: 60
    interval: 1
  ...
```

If you don't pass the value for `{{packetsize}}` while starting a task, the default one will be used.

Advanced templates

Yardstick makes it possible to use all the power of Jinja2 template syntax, including the mechanism of built-in functions. As an example, let us make up a task file that will do a block storage performance test. The input task file (`fio-template.yaml`) below uses the Jinja2 for-`endfor` construct to accomplish that:

```
#Test block sizes of 4KB, 8KB, 64KB, 1MB
#Test 5 workloads: read, write, randwrite, randread, rw
schema: "yardstick:task:0.1"

scenarios:
{% for bs in ['4k', '8k', '64k', '1024k' ] %}
  {% for rw in ['read', 'write', 'randwrite', 'randread', 'rw' ] %}
-
  type: Fio
  options:
    filename: /home/ec2-user/data.raw
    bs: {{bs}}
    rw: {{rw}}
    ramp_time: 10
  host: fio.demo
  runner:
    type: Duration
    duration: 60
    interval: 60

  {% endfor %}
{% endfor %}
context
  ...
```

YARDSTICK GLOSSARY

NFVI Network Function Virtualization Infrastructure

VNF Virtual Network Function

VTC Virtual Traffic Classifier

N

NFVI, 17

V

VNE, 17

VTC, 17