QTIP Userguide

Release draft (bccbfad)

OPNFV

August 16, 2016

CONTENTS

1 **Colorado 1.0** **1**

1.1 Overview 1

1.2 Qtip Introduction 1

CHAPTER

ONE

COLORADO 1.0

1.1 Overview

QTIP is an OPNFV Project.

QTIP aims to benchmark OPNFV platforms through a “Bottom up” approach, testing bare-metal components first.

The overall problem this project tries to solve is the general characterization of an OPNFV platform. It will focus on general performance questions that are common to the platform itself, or applicable to multiple OPNFV use cases. QTIP will provide the capability to quantify a platform’s performance behavior in a standardized, rigorous, and open way.

1.2 Qtip Introduction

This guide will serve as a first step to familiarize the user with how to run QTIP the first time when the user clones QTIP on to their host machine. In order to clone QTIP please follow the instructions in the installation.rst located in docs/userguide/installation.rst.

1.2.1 QTIP Directory structure:

The QTIP directory has been sectioned off into multiple folders to facilitate segmenting information into relevant categories. The folders that concern the end user are *test_cases/* and *test_list/*.

test_cases/:

This folder is used to store all the config files which are used to setup the environment prior to a test. This folder is further divided into opnfv pods which run QTIP. Inside each pod there are folders which contain the config files segmented based on test cases. Namely, these include, *Compute*, *Network* and *Storage*. The default folder is there for the end user who is interested in testing their infrastructure but arent part of a opnfv pod.

The structure of the directory for the user appears as follows

```
test_cases/default/compute
test_cases/default/network
test_cases/default/storage
```

The benchmarks that are part of the QTIP framework are listed under these folders. An example of the compute folder is shown below. Their naming convention is <BENCHMARK>_<VM/BM>.yaml

```
dhrystone_bm.yaml
dhrystone_vm.yaml
whetstone_vm.yaml
```

```
whetstone_bm.yaml
ssl_vm.yaml
ssl_bm.yaml
ramspeed_vm.yaml
ramspeed_bm.yaml
dpi_vm.yaml
dpi_bm.yaml
```

The above listed files are used to configure the environment. The VM/BM tag distinguishes between a test to be run on the Virtual Machine or the compute node itself, respectively.

test_list/:

This folder contains three files, namely *compute*, *network* and *storage*. These files list the benchmarks are to be run by the QTIP framework. Sample compute test file is shown below

```
dhrystone_vm.yaml
dhrystone_bm.yaml
whetstone_vm.yaml
ssl_bm.yaml
```

The compute file will now run all the benchmarks listed above one after another on the environment. *NOTE: Please ensure there are no blank lines in this file as that has been known to throw an exception.*

1.2.2 Preparing a config file for test:

We will be using dhrystone as a example to list out the changes that the user will need to do in order to run the benchmark.

Dhrystone on Compute Nodes:

QTIP framework can run benchmarks on the actual compute nodes as well. In order to run dhrystone on the compute nodes we will be editing the dhrystone_bm.yaml file.

```
Scenario:
  benchmark: dhrystone
  host: machine_1, machine_2
  server:
```

The *Scenario* field is used by to specify the name of the benchmark to run as done by *benchmark: dhrystone*. The *host* and *server* tag are not used for the compute benchmarks but are included here to help the user *IF* they wish to control the execution. By default both machine_1 and machine_2 will have dhrystone run on them in parallel but the user can change this so that machine_1 run dhrystone before machine_2. This will be elaborated in the *Context* tag.

```
Context:
  Host_Machines:
    machine_1:
      ip: 10.20.0.6
      pw:
      role: host
    machine_2:
      ip: 10.20.0.5
      pw:
      role: host

   Virtual_Machines:
```

The *Context* tag helps the user list the number of compute nodes they want to run dhrystone on. The user can list all the compute nodes under the *Host_Machines* tag. All the machines under test must be listed under the *Host_Machines* and naming it incrementally higher. The *ip:* tag is used to specify the IP of the particular compute node. The *pw:* tag can be left blank because QTIP uses its own key for ssh. In order to run dhrystone on one compute node at a time the user needs to edit the *role:* tag. *role: host* for machine_1 and *role: server* for machine_2 will allow for dhrystone to be run on machine_1 and then run on machine_2.

```
Test_Description:
  Test_category: "Compute"
  Benchmark: "dhrystone"
  Overview: >
      ''' This test will run the dhrystone benchmark in parallel  on
      machine_1 and machine_2.
```

The above field is purely for a description purpose to explain to the user the working of the test and is not fed to the framework.

1.2.3 Sample dhrystone_bm.yaml file:

```
Scenario:
  benchmark: dhrystone
  host: machine_1, machine_2
  server:

Context:
  Host_Machines:
    machine_1:
      ip: 10.20.0.6
      pw:
      role: host
    machine_2:
      ip: 10.20.0.5
      pw:
      role: host

  Virtual_Machines:


Test_Description:
  Test_category: "Compute"
  Benchmark: "dhrystone"
  Overview: >
      ''' This test will run the dhrystone benchmark in parallel  on
      machine_1 and machine_2.\n
```

Dhrystone on Virtual Machine:

To run dhrystone on the VMs we will be editing dhrystone_vm.yaml file. Snippets on the file are given below.

```
Scenario:
benchmark: dhrystone
host: virtualmachine_1, virtualmachine_2
server:
```

The *Scenario* field is used by to specify the name of the benchmark to run as done by *benchmark: dhrystone*. The *host* and *server* tag are not used for the compute benchmarks but are included here to help the user *IF* they wish to control the execution. By default both virtualmachine_1 and virtualmachine_2 will have dhrystone run on them in parallel but the user can change this so that virtualmachine_1 run dhrystone before virtualmachine_2. This will be elaborated in the *Context* tag.

```
Context:
  Host_Machines:

  Virtual_Machines:
    virtualmachine_1:
      availability_zone: compute1
      public_network: 'net04_ext'
      OS_image: QTIP_CentOS
      flavor: m1.large
      role: host
    virtualmachine_2:
      availability_zone: compute2
      public_network: 'net04_ext'
      OS_image: QTIP_CentOS
      flavor: m1.large
      role: host
```

The *Context* tag helps the user list the number of VMs and their characteristic. The user can list all the VMs they want to bring up under the *Virtual_Machines:* tag. In the above example we will be bringing up two VMs. One on Compute1 and the other on Compute2. The user can change this as desired *NOTE: Please ensure you have the necessary compute nodes before listing under the 'availability_zone:' tag*. The rest of the options do not need to be modified by the user.

Running dhrystone sequentially (Optional):

In order to run dhrystone on one VM at a time the user needs to edit the *role:* tag. *role: host* for virtualmachine_1 and *role: server* for virtualmachine_2 will allow for dhrystone to be run on virtualmachine_1 and then run on virtualmachine_2.

```
Test_Description:
  Test_category: "Compute"
  Benchmark: "dhrystone"
  Overview:
  This test will run the dhrystone benchmark in parallel on
  virtualmachine_1 and virtualmachine_2
```

The above field is purely for a decription purpose to explain to the user the working of the test and is not fed to the framework.

Running dhrystone with proxy (Optional):

In order to run the dhrystone on the hosts or vms which can only access the public network by proxy, the user needs to add the *Proxy_Environment* info in *Context* tag.

```
Context:
  Host_Machines:
    machine_1:
      ip: 10.20.0.29
      pw:
      role: host
    machine_2:
      ip: 10.20.0.30
      pw:
      role: host

  Virtual_Machines:

  Proxy_Environment:
    http_proxy: http://10.20.0.1:8118
    https_proxy: http://10.20.0.1:8118
    no_proxy: localhost,127.0.0.1,10.20.*,192.168.*
```

1.2.4 Sample dhrystone_vm.yaml file:

```
Scenario:
benchmark: dhrystone
host: virtualmachine_1, virtualmachine_2
server:

Context:
  Host_Machines:

  Virtual_Machines:
    virtualmachine_1:
      availability_zone: compute1
      public_network: 'net04_ext'
      OS_image: QTIP_CentOS
      flavor: m1.large
      role: host
    virtualmachine_2:
      availability_zone: compute2
      public_network: 'net04_ext'
      OS_image: QTIP_CentOS
      flavor: m1.large
      role: host

Test_Description:
  Test_category: "Compute"
  Benchmark: "dhrystone"
  Overview: >
  This test will run the dhrystone benchmark in parallel on
  machine_1 and machine_2.\n
```

1.2.5 Commands to run the Framework:

In order to start QTIP on the default lab please use the following commands (asssuming you have prepared the config files in the test_cases/default/ directory and listed the intended suite in the test_list/<RELEVANT-SUITE-FILE>):

First step is to export the necessary information to the environment.

```
source get_env_info.sh -n <INSTALLER_TYPE> -i <INSTALLER_IP>
```

for running qtip on an openstack deployed using FUEL with the Installer IP 10.20.0.2

```
source get_env_info.sh -n fuel -i 10.20.0.2
```

This will generate the *opnfv-creds.sh* file needed to use the python clients for keystone, glance, nova, and neutron.

```
source opnfv-creds.sh
```

Running QTIP on the using *default* as the pod name and for the *compute* suite

```
python qtip.py -l default -f compute
```

Running QTIP on the using *default* as the pod name and for the *network* suite

```
python qtip.py -l default -f network
```

Running QTIP on the using *default* as the pod name and for the *storage* suite

```
python qtip.py -l default -f network
```

1.2.6 Results:

QTIP generates results in the *results/* directory are listed down under the particularly benchmark name. So all the results for dhrystone would be listed and time stamped.