

Provisioning Using Opendaylight OVSDb Into Openstack: Experiments

Alexandru Eftimie,
University POLITEHNICA of Bucharest, Bucharest,
Romania
Department of Telecommunication, University
"Politehnica" of Bucharest, , Romania
E-mail: alexandru.eftimie@gmail.com

Eugen Borcoci
University POLITEHNICA of Bucharest, Bucharest,
Romania
Department of Telecommunication, University
"Politehnica" of Bucharest, , Romania
E-mail: eugen.borcoci@elcom.pub.ro

Abstract – Network function virtualization (NFV) and Software Defined Networking (SDN) are complementary technologies that support flexible development or virtual machines in various environments, e.g., multi-tenant / multi-domain. While SDN separates the architectural control plane, versus data plane, NFV implements a lot of functions (that traditionally have been performed by dedicated boxes), by software – using virtualized network functions (VNF). There are still open research issues in both SDN and NFV, especially related to their cooperation and integration. This paper presents an experiment of deploying networks and virtual machines (VMs) using Openstack and Open vSwitch Database Management Protocol (OVSDb) from Opendaylight project. The study is oriented towards implementation aspects. Its main objective is to deploy an Openstack controller, an Opendaylight controller and two compute nodes and to create on top of existing infrastructure several networks and illustrate how this is automatically achieved and how overlay networks can coexist. The paper describes step by step how to configure controllers, how connectivity is achieved and how OpenFlow is used to forward packets.

Keywords-Openstack; Network Function Virtualization; Software Defined Network; OpenFlow.

I. INTRODUCTION

In traditional networking architecture, the IP datagrams are carried and processed by network nodes that are individual boxes, each performing a specific function, such as forwarding, switching, filtering, firewall. These network boxes bring costs, capital expenditure (CAPEX) and operating expenses (OPEX), and there is no flexibility in using them as the network is growing and, most important, changing.

Software Defined Networking (SDN) separates network control and data forwarding functions leading to centralized and programmable network control. SDN architecture has several main components such as: data plane consisting in network resources for forwarding traffic; control plane implemented as SDN controller, which manages the network resources; network applications plane. The interface between the forwarding

and the control plane is the “southbound” interface, and the interface between the control plane and applications is the “northbound” [1].

Network function virtualization (NFV) aims to implement by software many functions, that traditionally have been implemented as expensive hardware-software combinations. Recent standards define the NFV architecture and also how to implement different virtualized network functions (VNF) in a virtual environment – to replace the traditional dedicated boxes which performed individual functions. The SDN and NFV are complementary technologies, usable independently or in cooperation. While NFV replaces hardware network elements, SDN deals with replacement of network protocols, bringing centralized control. [2]

By decoupling the two planes (SDN) and by using the function virtualization (NFV) many borders of traditional networks can be overcome. All functions that are currently delivered by hardware boxes will be implemented in software, the deployments can be done automatically, on demand or by reacting to network changes.

This paper is organized in three sections. Section number II is an introduction to the main technologies used and is describing the relevance of NFV and SDN. Section III will go through all the implementation steps and details demonstrating the cooperation between the technologies and describing the results.

II. NETWORK FUNCTION VIRTUALIZATION (NFV) and SOFTWARE DEFINED NETWORK (SDN) INTEGRATION

This section will summarize some aspects of Network Function Virtualization and Software Defined Network and introduce the SDN-NFV approach.

A. Network Function Virtualization

In traditional networks that do not yet benefit from virtualization, network features are implemented as a combination of software and hardware that are specific to a manufacturer; these are called network nodes. Virtualization of network functions is a step forward in the telecommunications environment by introducing

differences in the way services are delivered compared to the current practice. These differences can be described as the following:

- Decoupling hardware from software;
- Flexible implementation of network feature;
- Dynamic operations.

Implementation, control and management of network functions are required in the context of NFV-enabled network nodes for various optimization purposes. Thus, many challenges relate to the algorithm and design of the system in terms of implementation of functions. One of these challenges is the automatic provision of network and processing resources based on the use of the underlying resources involved. A similar and probably the most important challenge is the placement and automatic allocation of VNFs, because their placement and allocation significantly influence service performance. Both automatic supply and placement require a global view of resources and a unified control and optimization system with various optimization engines running in it [2].

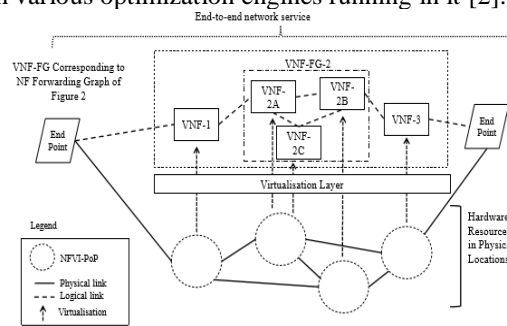


Figure 1: Example of end-to-end network service using VNFs and forwarding graph [2]

The end-to-end network services can be composed of several VNFs, organized in forwarding graphs (VNF-FG), as depicted in Figure 1. Terminals and network functions are nodes and correspond to equipment, applications, or even physical servers.

B. Software Defined Network

By separating control plane from data plane, the network switches become mainly forwarding devices. However, the SDN flow concept (and the flow tables installed in the SDN switches by the controller), allows a large range of processing actions to be executed upon the packets of a flow matching the flow tables. So, an SDN switch can be more powerful than a traditional router. A simplified view of SDN architecture is shown in Figure 2.




Figure 2: SDN Architecture [3]

The OpenDaylight project [4] is an open source SDN platform that uses open protocols to provide centralized control and monitor of network devices. Like many other SDN controllers, OpenDaylight supports OpenFlow, offering network ready solutions for installation as part of the platform.

The core of the OpenDaylight platform is the Model-Drive Service Abstraction Layer (MD-SAL). In OpenDaylight, network-based devices and network applications are represented as objects or models; SAL is a mechanism for data exchange and adaptability between YANG models, representing network devices and applications. YANG models provide generalized descriptions of capabilities of a device or application without the need to know specific details of their implementation [5] [4].

C. Integration

In the SDN-NFV approach the network functions are implemented as software modules, running on virtual machines with the control of a hypervisor, which allows the flexibility of supplying computing and network resources. Thus, because the computational capacity can be increased when needed, there is no need for over-provisioning. On the other hand, service chaining in SD-NFV benefits from improvement. For geographically spread networks, updating devices requires a high cost. Additionally, errors may occur in update operations, and reconfiguration may lead to disruption of the entire network.

OpenStack and OpenDaylight Integration




Figure 3: Openstack – OpenDaylight integration [6]

However, with SD-NFV, service providers can create new connections without radically changing hardware.

With intelligent service linkages, the complexity of resource delivery is significantly reduced. The SD-NFV architecture is still in the research phase; one possible integration is represented in Figure 3; a single control and orchestration framework is required to integrate the SDN controller, routing elements and virtual network functions. Moreover, due to the functionality dynamic feature and the provision of resources, this SD-NFV framework must provide coordinated control status [1] [7]. The Modular Layer 2 (ML2) plugin is a framework allowing OpenStack Networking to simultaneously utilize the variety of layer 2 networking technologies found in real-world data centers.

III. OVSDB WITH OPENSTACK USE CASE

Openstack offers open APIs to support a wide range of applications and infrastructures, including Neutron API and Neutron / Multi-Layer 2 (ML-2) for networking. Neutron offers a "low-level" interface and was not designed to manage the data center substrates. The ML-2 Neutron has been designed to expose data center switch capabilities, but there is currently a limitation to some virtual and physical switches. (Opendaylight - Cloud and NFV)

Opendaylight is an open source framework for migrating to a SDN network architecture [8].

This section will cover all the steps done to configure and deploy an Openstack environment using OVSDB project from Opendaylight. This environment will be used to deploy network and instances on various compute node and to show how tunnels are automatically configured and how OpenFlow is used for forwarding the traffic [6].

A. Architecture and IP addressing schema

For this experiment, a Fedora 32-bit image was used, with an average resource allocation as there will be several machines in place. For the experiment creating and managing the network using an Openstack controller and ODL controller through OpenFlow, 3 virtual machines were chosen. A virtual machine will act as a controller and two will be compute nodes. The used hypervisor is the Oracle Vm Virtual Box that was configured with two private networks, VirtualBox Host-Only Ethernet Adapter and VirtualBox Host-Only Ethernet Adapter 2, that are allocating IP addresses via DHCP from 192.168.56.0/24 and 192.168.57.0/24. The addressing scheme is described in Table 1 and the interconnect and each machine role is depicted in Figure 4.

After booting the VM's the IP addressing is the following along with the roles inside the Openstack/ODL architecture:

TABLE I ADDRESSING SCHEME OF THE THREE VMs

VM	Role	Interface p7p1	Interface p8p1
Fedora 1	Controller node	192.168.56.117	192.168.57.113
Fedora 2	Compute node	192.168.56.118	192.168.57.114
Fedora 3	Compute node	192.168.56.119	192.168.57.115




Figure 4: Roles of VMs

To prepare the setup for deployment, there are several steps required:

- a) Start the OVS service - although this service should start automatically when running the devstack configuration script of Openstack, we can start the service using the following command:

```
[fedora@fedora1 ~]$ sudo /sbin/service
opendaylight start
Redirecting to /bin/systemctl start
opendaylight.service
```

- b) Configure the /etc/hosts file to reflect the host and IP addressing of the lab. This file will look the same on all three machines:

```
[fedora@fedora1 ~]$ sudo vi /etc/hosts
127.0.0.1          fedora1          localhost
localhost.localdomain  localhost4
localhost4.localdomain4
192.168.56.117 fedora1
192.168.56.118 fedora2
192.168.56.119 fedora3
::1               localhost localhost.localdomain
localhost6 localhost6.localdomain6
```

- c) Change the hostname on the VMs and reboot for the changes to take effect:

```
[fedora@fedora1 ~]$ cat /etc/hostname
[fedora@fedora1 ~]$ sudo hostname -b
fedora2
[fedora@fedora1 ~]$ sudo shutdown -r now
```

B. Starting the Opendaylight controller on the Openstack controller

Configuration must be checked to be sure that the OpenFlow version used is 1.3 as all the scripts will ask for this version:

```
[fedora@fedora1 ~]$ cd opendaylight/
[fedora@fedora1 opendaylight]$ grep
ovsdb.of.version configuration/config.ini
ovsdb.of.version=1.3
```

Next action required is to launch the Opendaylight controller, using:

```
./run.sh -XX:MaxPermSize=384m -virt ovsdb
-of13
```

Once started the deployment is done automatically and it will be finished when the FlowConfiv provider, GroupConfig provider, MeterConfig Provider and the Statistics Provider are started

```
2019-05-07 12:21:34.153 CEST [pool-2-
thread-5] INFO
o.o.controller.frm.flow.FlowProvider - Flow
Config Provider started.
2019-05-07 12:21:34.159 CEST [pool-2-
thread-5] INFO
o.o.c.frm.group.GroupProvider - Group
Config Provider started.
2019-05-07 12:21:34.226 CEST [pool-2-
thread-5] INFO
o.o.c.frm.meter.MeterProvider - Meter
Config Provider started.
2019-05-07 12:21:34.247 CEST [pool-2-
thread-4] INFO
o.o.c.m.s.manager.StatisticsProvider -
Statistics Provider started.
The controller will start listen on TCP port 6633:
[fedora@fedora1 ~]$ lsof -iTCP | grep 66
java 1223 fedora 41u IPv6 20211
0t0 TCP *:6633 (LISTEN)
java 1223 fedora 57u IPv6 20216
0t0 TCP *:6653 (LISTEN)
```

C. Configure Openstack controller

Openstack must be configured and the controller must be started using Devstack tool. It is required to keep the ODL controller started and a new SSH connection is made to Fedora 1. On Fedora 2 and 3 it can be done from same command line.

- a) Edit the local.conf file according to our IP addressing schema – I will copy here only the most important elements that need configuration:

```
[[local|localrc]]
LOGFILE=stack.sh.log
```

```
SCREEN_LOGDIR=/opt/stack/data/log
LOG_COLOR=False
OFFLINE=True
#RECLONE=yes
HOST_IP=<IP-ADDRESS-OPENSTACK-CONTROLLER>
HOST_NAME=fedora1
SERVICE_HOST_NAME=${HOST_NAME}
SERVICE_HOST=<IP-ADDRESS-OF-OPENSTACK-
CONTROLLER>
url=http://<IP-ADDRESS-OF-ODL-
CONTROLLER>:8080/controller/nb/v2/neutron
```

```
[fedora@fedora1 devstack]$ grep 192.168.56
local.conf
HOST_IP=192.168.56.117
SERVICE_HOST=192.168.56.117
url=http://192.168.56.117:8080/controller/
nb/v2/neutron
```

SERVICE_HOST and the IP address from the url will always point to the controller, all other elements will be configured with local information. Having these changes made, Openstack can be started by running the stack.sh script and the installation is done automatically. During the installation it can be observed that in the ODL command line there are several messages showing the communication between ODL and Openstack [9]:

```
2019-05-07 12:35:54 Starting Neutron
2019-05-07 12:35:56 Waiting for Neutron to
start...
2019-05-07 12:36:04 {"versions":
[{"status": "CURRENT", "id": "v2.0",
"links": [{"href":
"http://192.168.56.117:9696/v2.0", "rel":
"self"}]}]}Added interface d0ced68a-2350-
4516-9908-a073fe208af2 to router 9a89b604-
750f-4f7b-a440-ce22e78321e1.
```

```
osgi> 2019-05-07 12:35:59.978 CEST [http-
bio-8080-exec-1] INFO
o.o.c.u.internal.UserManager - Local
Authentication Succeeded for User: "admin"
2019-05-07 12:35:59.980 CEST [http-bio-
8080-exec-1] INFO
o.o.c.u.internal.UserManager - User "admin"
authorized for the following role(s):
[Network-Admin]
```

```
Open vSwitch after the compute nodes are stacked:
[fedora@fedora1 devstack]$ sudo ovs-vsctl
show
3cc9dac3-9fa6-4c69-acc1-a2d463396fc8
Manager "tcp:192.168.56.117:6640"
is_connected: true
```

```

Bridge br-int
  Controller
"tcp:192.168.56.117:6633"
  is_connected: true
  fail_mode: secure
  Port "vxlan-192.168.56.118"
  Interface "vxlan-192.168.56.118"
  type: vxlan
  options: {key=flow,
local_ip="192.168.56.117",
remote_ip="192.168.56.118"}
    
```

One thing to note is that the manager is configured automatically, and it always points to the ODL controller. The connection is done to OVSDB socket 192.168.56.117:6640. Also, the br-int bridge is created on all 3 instances - the controller is the same machine as the IP address 192.168.56.117, but the connection is made on port 6633 that indicates the **connection to OpenFlow**. During tests it has been noticed that it is mandatory to have the status of "is_connected" to be true. If it is not present, OVS is configured, but the connection is not available (no errors are being throwned, needs manual check). Same procedure needs to be done on Fedora 2 and Fedora 3.

D. Provisioning

Before starting to provision the infrastructure it must be checked if there are three hypervisors registered with Nova:

- Populate the proper Keystone credentials for service client commands using the openrc file:

```
[fedora@fedora1 devstack]$ ./openrc admin admin
```

- Check hypervisor's list:

```
[fedora@fedora1 devstack]$ nova hypervisor-list
+-----+-----+
| ID | Hypervisor hostname |
+-----+-----+
| 1 | fedora1              |
| 2 | fedora2              |
| 3 | fedora3              |
+-----+-----+
```

- Run the add.ingage.sh script - this is required because the virtual machine we're working on is a Fedora 32-bit architecture and default Cirrus image is not 32 bit, so we add an image with x386 architecture in Glance (cirros-0.3.1-i386)

```
[fedora@fedora1 devstack]$ cat ./addimage.sh
```

```
[fedora@fedora1 devstack]$ ./addimage.sh
Added new image with ID: a335d33d-fee1-41be-8be9-458cd3f7e309
/home/fedora/devstack
```

- On this infrastructure 6 overlay network will be build, using GRE and VxLAN encapsulation:

```
neutron net-create gre1 --tenant_id $(keystone tenant-list | grep '\sadmin' | awk '{print $2}') --provider:network_type gre --provider:segmentation_id 1300
neutron subnet-create gre1 10.100.1.0/24 --name gre1
```

A summary of those networks will look like this:

id	name
45cf5f88-1d60-46da-a651-a3df32bc217c	gre1
4e5c0d59-e903-47a1-8bbc-3691351eb4ce	vxlan-net1
7e634fbe-ebcf-47a3-80af-7054ad4fb4fb	gre2
9574dfb9-4e52-4243-8b5e-10f1b44bfab1	private
b980e7ba-c495-44ae-8dc8-7066ac7fd941	gre3
d56a4ac9-34f0-45d2-be92-2d9d7ae83153	public
ef07fbf4-1b98-42ff-92c4-87dcb79dcaff	vxlan-net3
f2b32863-39d1-459c-97c5-7f9212e618e0	vxlan-net2

- Starting instances on compute nodes:

To create VM on the command line, the following commands were used, specifying which networks are to be attached. There is also the option to use the "availability zone" option to specify which compute nodes will host the VMs

```
nova boot --flavor m1.tiny --image $(nova image-list | grep $IMAGE\s' | awk '{print $2}') --nic net-id=$(neutron net-list | grep vxlan-net1 | awk '{print $2}') vxlan-host1 --availability_zone=nova:fedora2
```

ID	Name	Status
e02df395-d234-4995-bf69-2db77cb7d3ec	admin-private1	ACTIVE
278619f0-9a5d-4f0d-a336-16d0d282c4a6	gre-host1	ACTIVE
30e30c81-c7a0-46d7-869f-4229ef5e0690	gre-host2	ACTIVE
5c808b8a-5c80-4c39-9464-4b25c2fe4c34	vxlan-host1	ACTIVE
6d26fc2a-bb45-4eef-903f-7f115ac4198a	vxlan-host2	ACTIVE

One can observe on Fedora 2 that GRE and VxLAN tunnels have been automatically created between neighboring switches. We get a full mesh between VM in this way and control their creation by specifying the location for each instance.

```
[fedora@fedora3 devstack]$ sudo ovs-vsctl show
```

```

3cc9dac3-9fa6-4c69-acc1-a2d463396fc8
  Manager "tcp:192.168.56.117:6640"
  is_connected: true
  Bridge br-int
  Controller
"tcp:192.168.56.117:6633"
  is_connected: true
  Port br-int
    type: internal
  Interface "tap86dfa951-8b"
  Port "vxlan-192.168.56.117"
  Interface "vxlan-192.168.56.117"
    type: vxlan
    options: {key=flow,
local_ip="192.168.56.119",
remote_ip="192.168.56.117"}
  Port "gre-192.168.56.118"
  Interface "gre-192.168.56.118"
    type: gre
    options: {key=flow,
local_ip="192.168.56.119",
remote_ip="192.168.56.118"}
  Port "tapc42b0d97-43"
  Interface "tapc42b0d97-43"
  Port "vxlan-192.168.56.118"
  Interface "vxlan-192.168.56.118"
    type: vxlan
    options: {key=flow,
local_ip="192.168.56.119",
remote_ip="192.168.56.118"}
  Port "gre-192.168.56.117"
  Interface "gre-192.168.56.117"
    type: gre
    options: {key=flow,
local_ip="192.168.56.119",
remote_ip="192.168.56.117"}

```

The network topology can also be checked via Openstack Horizon at <http://192.168.56.117>, using default credentials (user: admin, password: admin):




Figure 5: Network topology from Openstack web interface

In Figure 5 it can be seen the network topology obtained, each network being represented in different color and having corresponding hosts attached.

E. End-to-end connectivity

To understand how traffic is routed to such an infrastructure, OpenFlow inputs must be analyzed. For the present scenario, there are 3 OpenFlow tables: Table 0, Table 10, and Table 20. Table 0 is the default table. For each tunnel created, there is an OpenFlow port present in this table, and we can see that all these tunnels are finished in this table - the full mesh network. Mapping between tunnels and OpenFlow ports in this table is done using "tun_id" as a key.

In addition to provisioning, which involves networking, sub-networks, encapsulation configuration, and the launch of instances on the available infrastructure that has been presented so far, functional verification requires end-to-end connectivity between the controller and the created instance. Using the "nova list" command, we visualize the created instances and choose one of these to check for connectivity. The result of this command tells us for each instance the network to which it is attached as well as its IP address.

To test connectivity, use the ping utility, but for this packets must be sent using the interface that is connected to the same network. As shown in Figure 6 the identifier of the dhcp server dealing with addressing in the gre1 10.100.1.0/24 network was used, this being the network attached to the chosen instance.

```

[fedora@fedora1 devstack]$ sudo ip netns exec
qdhcp-45cf5f88-1d60-46da-a651-a3df32bc217c ping
10.100.1.2
PING 10.100.1.2 (10.100.1.2) 56(84) bytes of data:
64 bytes from 10.100.1.2: icmp_seq=1 ttl=64
time=22.3 ms
64 bytes from 10.100.1.2: icmp_seq=8 ttl=64
time=12.0 ms
^C
--- 10.100.1.2 ping statistics ---
8 packets transmitted, 8 received, 0% packet loss,
time 7008ms
rtt min/avg/max/mdev = 0.392/7.810/22.338/7.077 ms

```

Figure 6 PING experiment - SUCCESS

To understand how traffic flows there were selected from the OpenFlow tables all entries used to forward the traffic:

<p>Tabela 0: cookie=0x0, duration=7472.684s, table=0, n_packets=120, n_bytes=11850, send_flow_rem in_port=8,dl_src=fa:16:3e:f5:2c:a0</p>
<p>Tabela 10: cookie=0x0, duration=921.617s, table=10, n_packets=210, n_bytes=20446, send_flow_rem tun_id=0x514,dl_dst=fa:16:3e:e0:2e:b 8 actions=output:11,goto_table:20</p>
<p>Tabela 20: cookie=0x0, duration=7587.297s, table=20, n_packets=238, n_bytes=22652, send_flow_rem tun_id=0x514,dl_dst=fa:16:3e:f5:2c:a 0 actions=output:8</p>

Figure 7: OpenFlow tables

In Figure 7 are represented a couple of entries from tables 0, 10 and 20. Those tables are used to forward packets between the tables or directly to the physical port.

IV. CONCLUSION

This paper presented an experiment using Openstack and Opendaylight framework used to provision and manage an infrastructure. As a proof of concept connectivity end to end between newly deployed VM and controller was successful and RTT achieved comparing to the one between two Fedora machines shows that the performance is not the same, but similar – the amount of resource allocated per machine was limited and the connectivity between instance and controller requires additional encapsulation. The paper can help designers to develop and implement systems based on cooperation between NFV and SDN, applicable to future IT developments and for the upcoming 5G technology. It is very important to choose wisely the operating system as there are unstable versions for such scenarios. The small setup implemented in this paper can be scaled to a data center to deploy and manage a larger number of instances and traffic.

As future work several experiments will be conducted, including building a network function chaining using similar deployments and compare from automation, resource consumption and stability perspective.

REFERENCES

- [1] Y. Li and M. Chen, "Software-Defined Network Function Virtualization: A Survey," *IEEE Xplore Digital Library*, pp. 2542 - 2553, 9 December 2015.
- [2] ETSI GS NFV 002 V1.1.1 (2013-10) , "Network Functions Virtualisation (NFV); Architectural Framework," *ETSI*, 2013.
- [3] K. Diego , M. V. R. Fernando , E. V. Paulo, E. R. Christian, A. M Siamak and U. Steve, "Software-Defined Networking: A Comprehensive Survey," *IEEE Xplore Digital Library*, pp. 14-76, 19 Decembrie 2014.
- [4] Opendaylight community, "Platform Overview," [Online]. Available: <https://www.opendaylight.org/what-we-do/odl-platform-overview>. [Accessed 15 December 2019].
- [5] RFC 6020, "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)," Oct 2010. [Online]. Available: <https://tools.ietf.org/html/rfc6020>. [Accessed Jan 2020].
- [6] Opendaylight community, "ODL Cloud and NFV," [Online]. Available: <https://www.opendaylight.org/use-cases-and-users/by-function/cloud-and-nfv>. [Accessed 15 December 2019].
- [7] Q. Duan, N. Ansari and M. Toy, "Software-Defined Network Virtualization: An Architectural Framework for Integrating SDN and NFV for Service Provisioning in Future Networks," *IEEE Network*, 2016.
- [8] Openstack community, "Openstack," 2018. [Online]. Available: <https://docs.openstack.org/security-guide/introduction/introduction-to-openstack.html>. [Accessed 15 December 2019].
- [9] B. Salisbury, "Networkstatic," [Online]. Available: <http://networkstatic.net/opendaylight-openstack-integration-devstack-fedora-20/#!prettyPhoto>. [Accessed 15 December 2019].