CHINA
OpenStack Days

基于Tacker+Networking-SFC的服务链编排方案实践

任亮
山石网科资深产品技术专家

1 NFV架构介绍

2 OPNFV的实践

3 Tacker+Networking-SFC的实践

4 遗留问题

ETSI定义的NFV框架

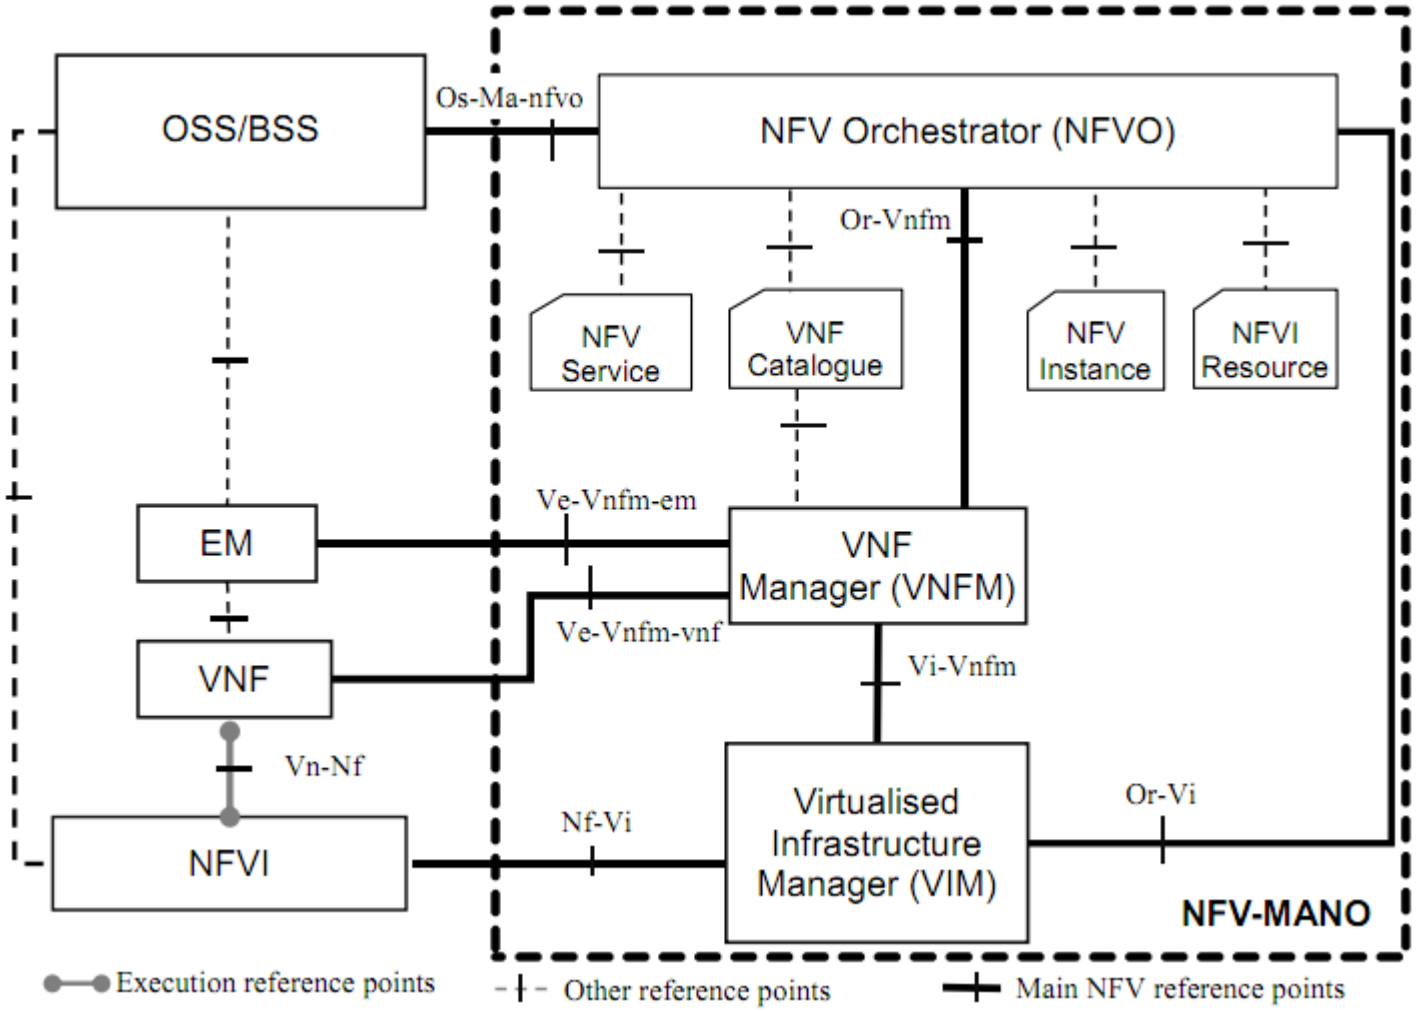

OpenStack的MANO组件—Tacker

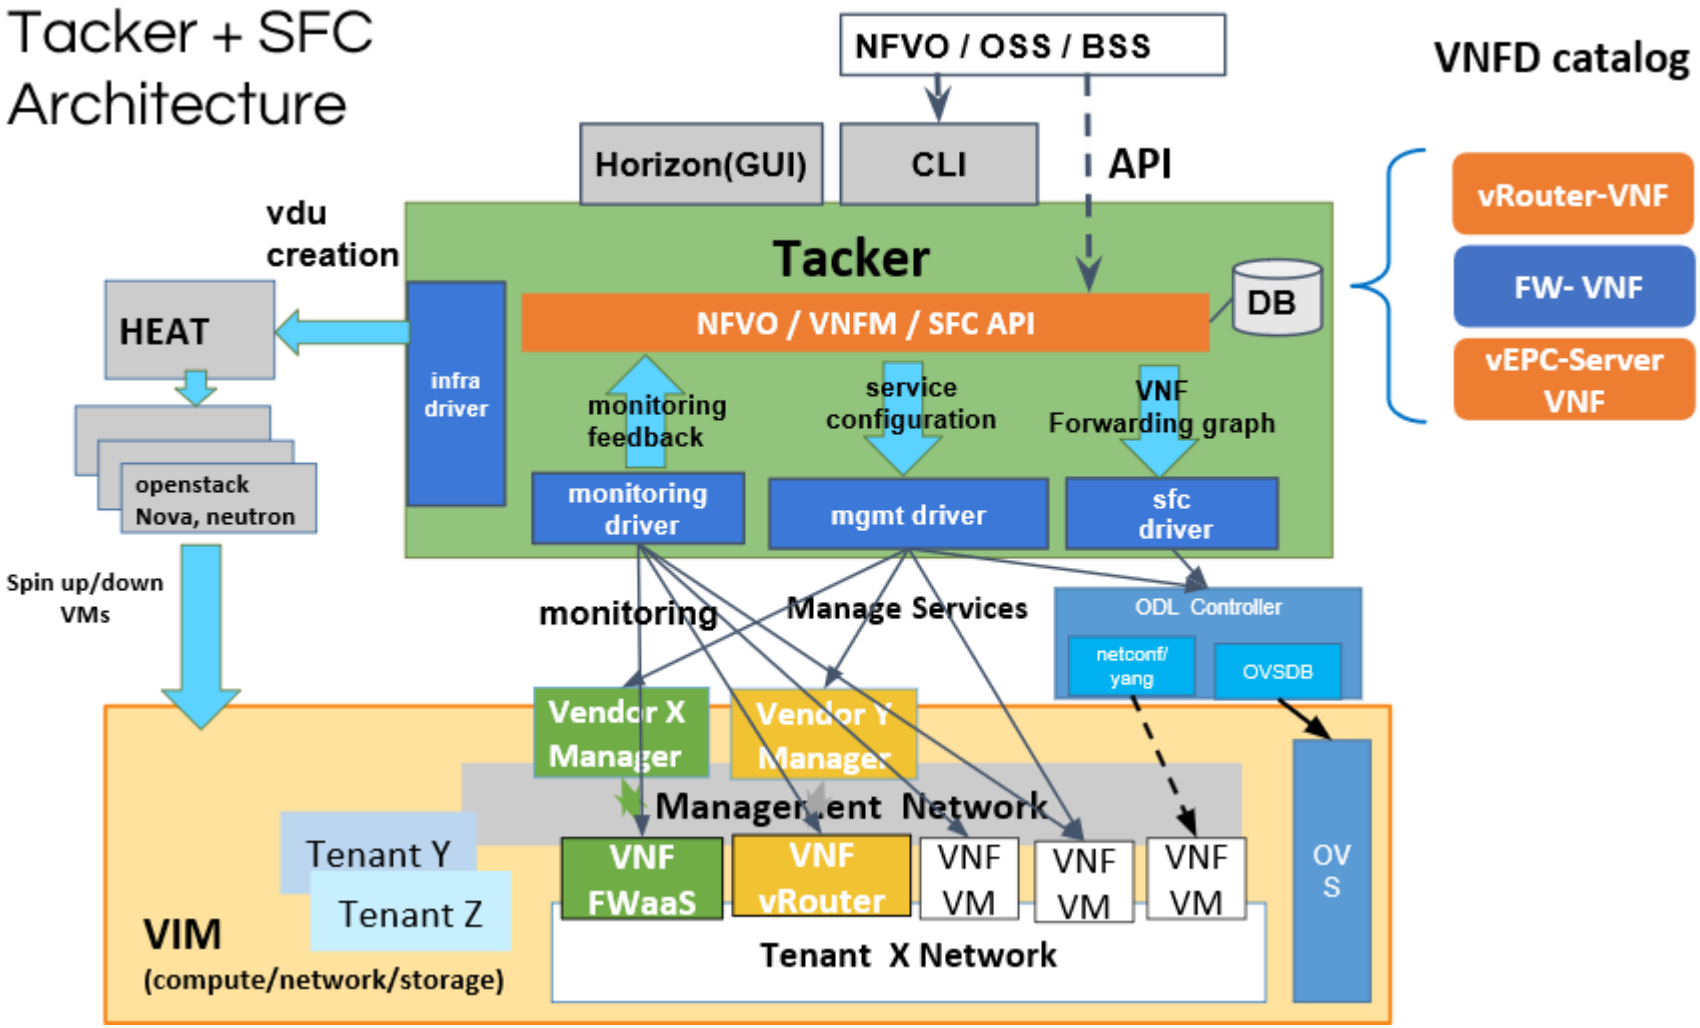

OPNFV的框架

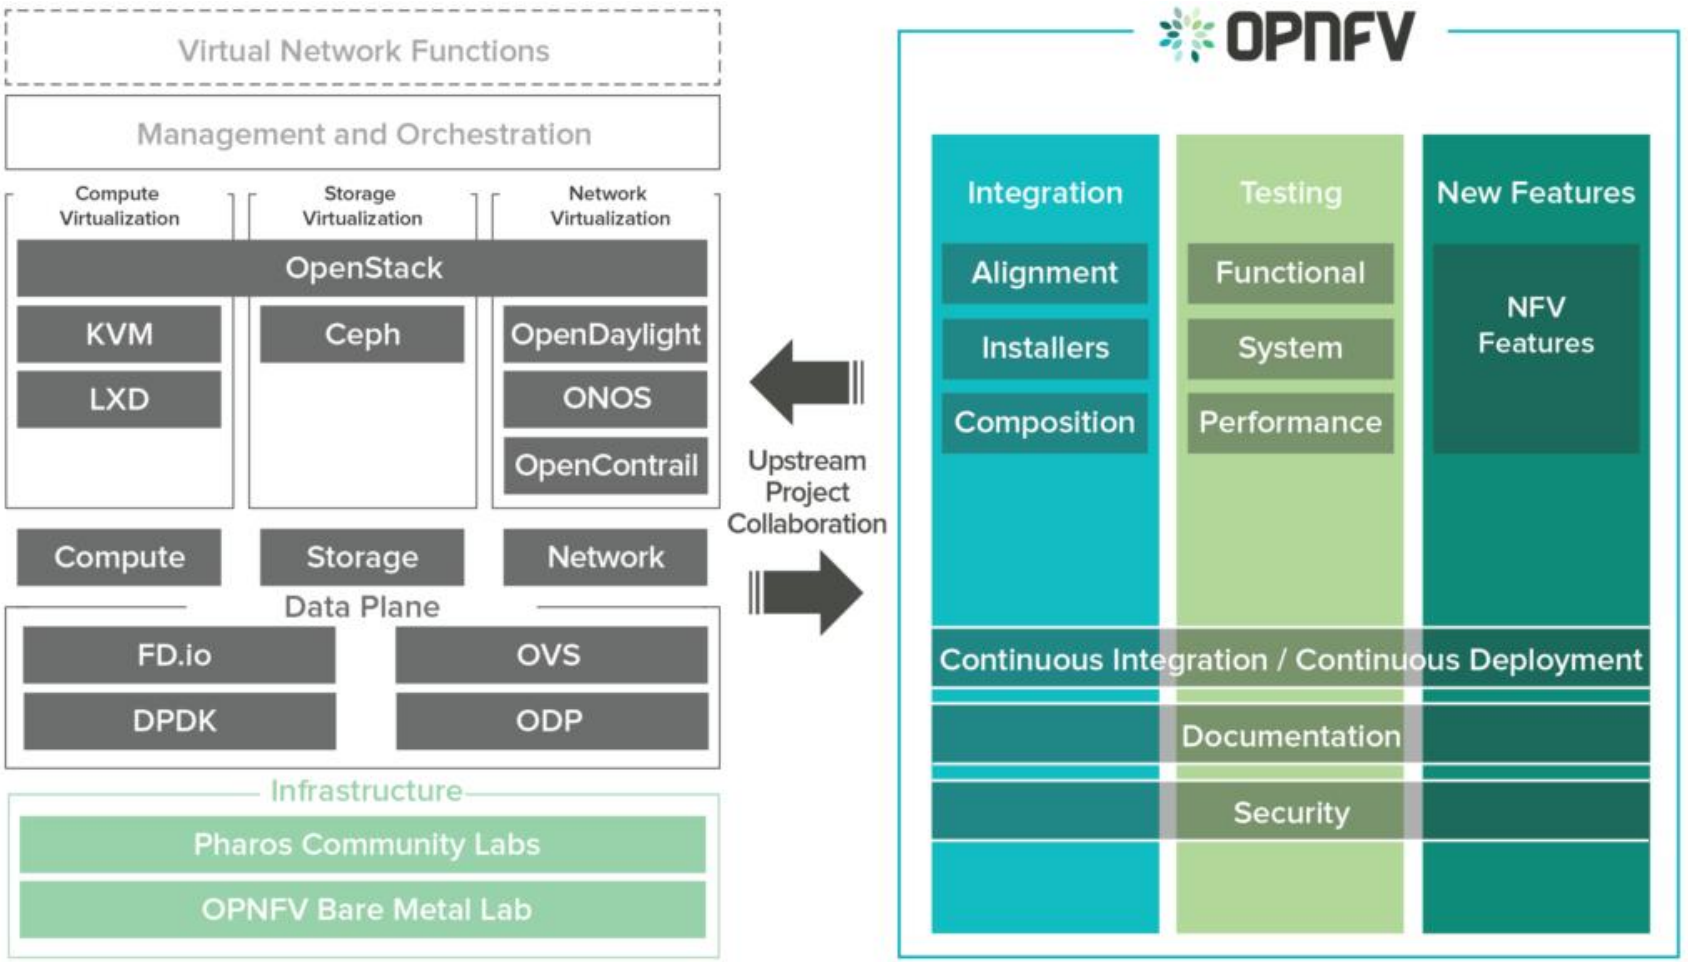
Virtual Network Functions
Management and Orchestration
Compute Virtualization
Storage Virtualization
Network Virtualization
OpenStack
KVM
LXD
Ceph
OpenDaylight
ONOS
OpenContrail
Compute
Storage
Network
Data Plane
FD.io
DPDK
OVS
ODP
Infrastructure
Pharos Community Labs
OPNFV Bare Metal Lab
Upstream Project Collaboration
OPNFV
Integration
Alignment
Installers
Composition
Testing
Functional
System
Performance
New Features
NFV Features
Continuous Integration / Continuous Deployment
Documentation
Security

OPNFV的SFC方案

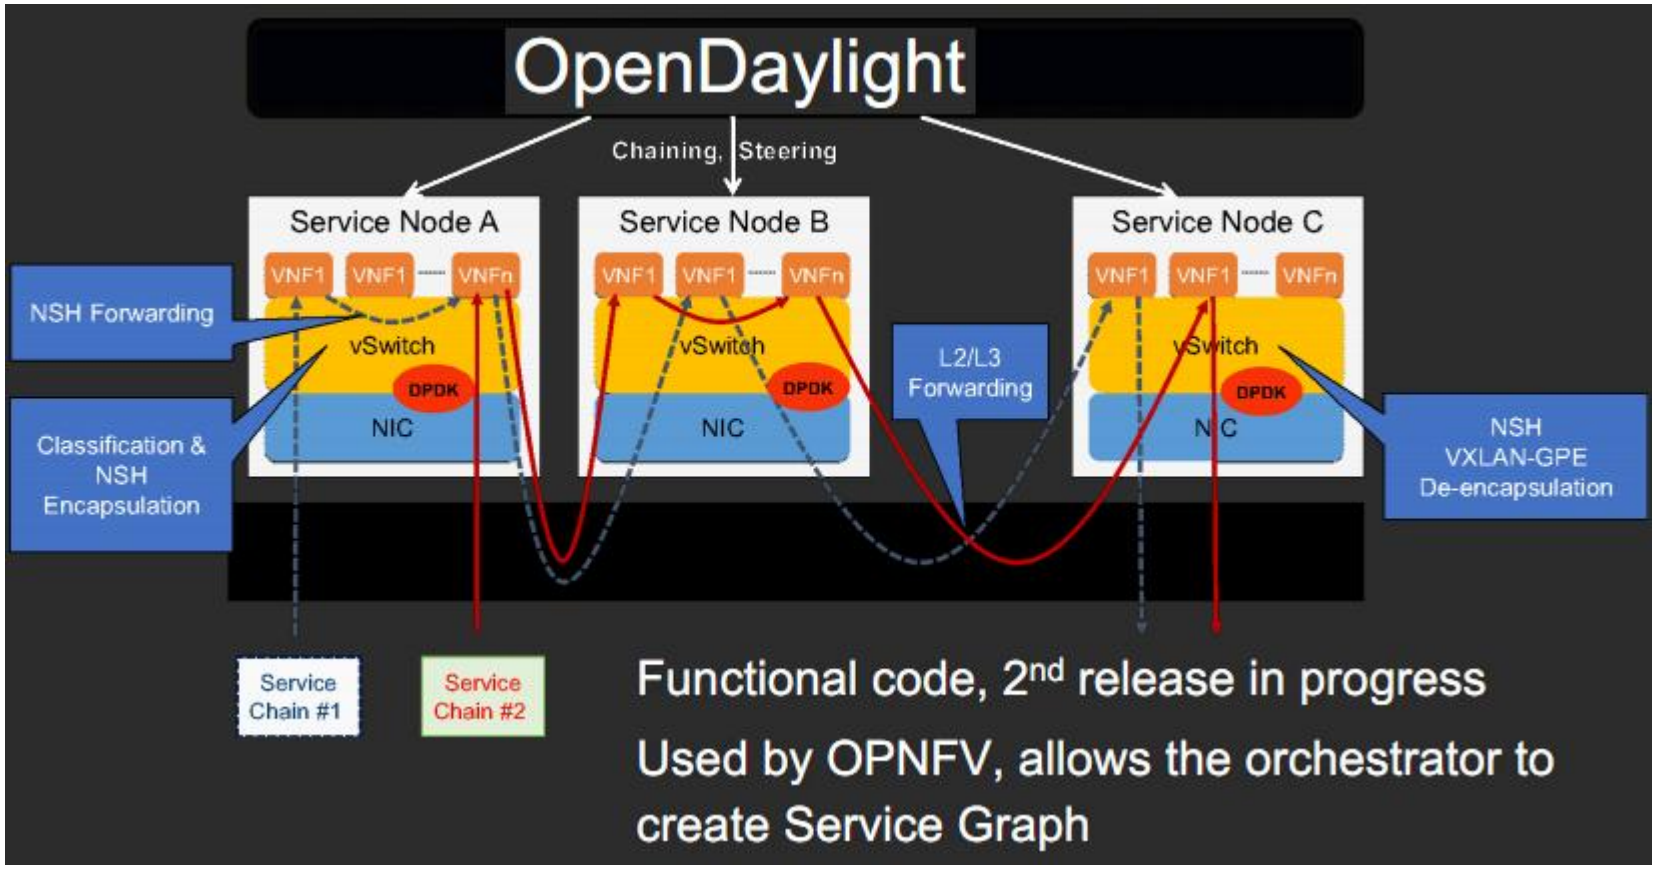

1 NFV架构介绍

2 OPNFV的实践

3 Tacker+Networking-SFC的实践

4 遗留问题

OPNFV的NFV实践—Danube 2.0

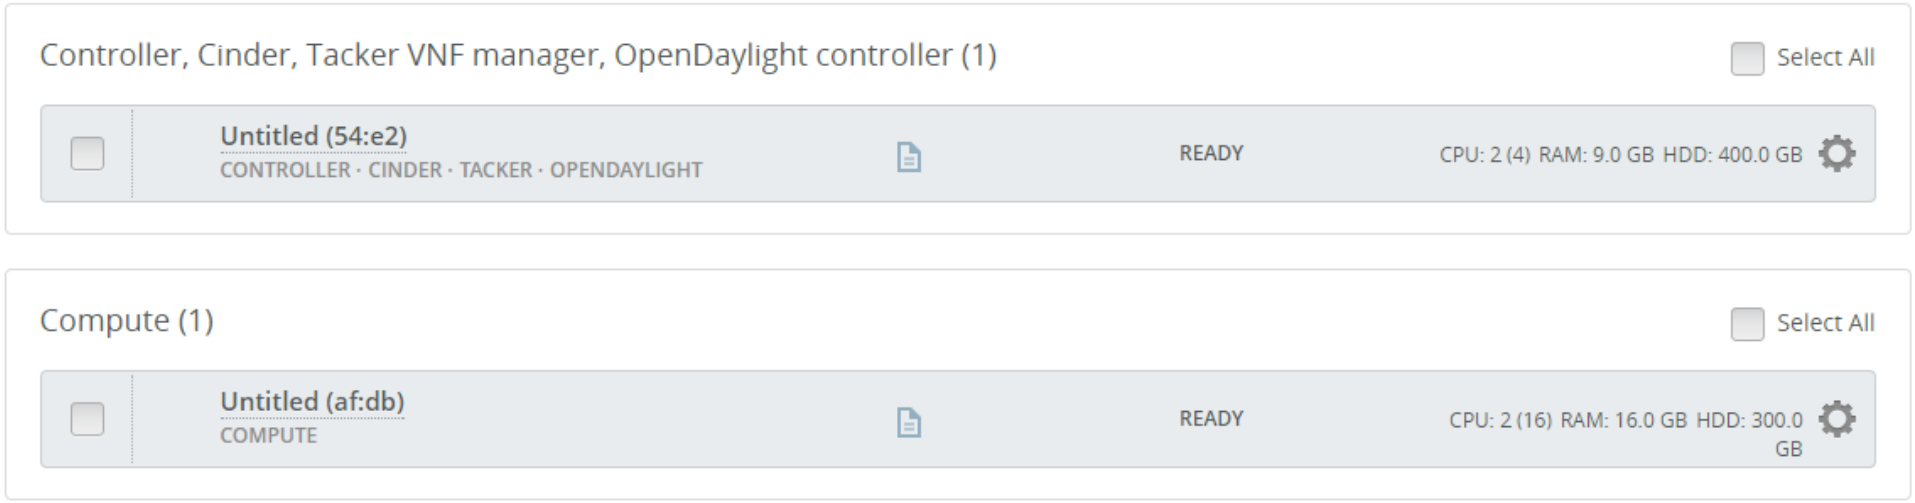

http://artifacts.opnfv.org/fuel/danube/opnfv-danube.2.0.iso

OPNFV实践—Plugin

Install Open vSwitch with NSH

Plugin version:	1.0.0
Description:	This plugin provides to deploy Open vSwitch with NSH
Authors:	ruijing.guo@intel.com
Licenses:	Apache License Version 2.0
Releases:	Ubuntu: newton-10.0

OpenDaylight plugin

Plugin version:	1.0.0
Description:	This plugin provides OpenDaylight as a backend for neutron.
Authors:	Michal Skalski, Nikolas Hermanns
Licenses:	Apache License Version 2.0
Releases:	Ubuntu: newton-10.0

Tacker VNF manager

Plugin version:	1.0.0
Description:	Tacker VNF manager
Authors:	Michal Skalski, Mirantis'Priya, Cavium
Licenses:	Apache License Version 2.0
Releases:	Ubuntu: newton-10.0

OPNFV的实践—Plugin

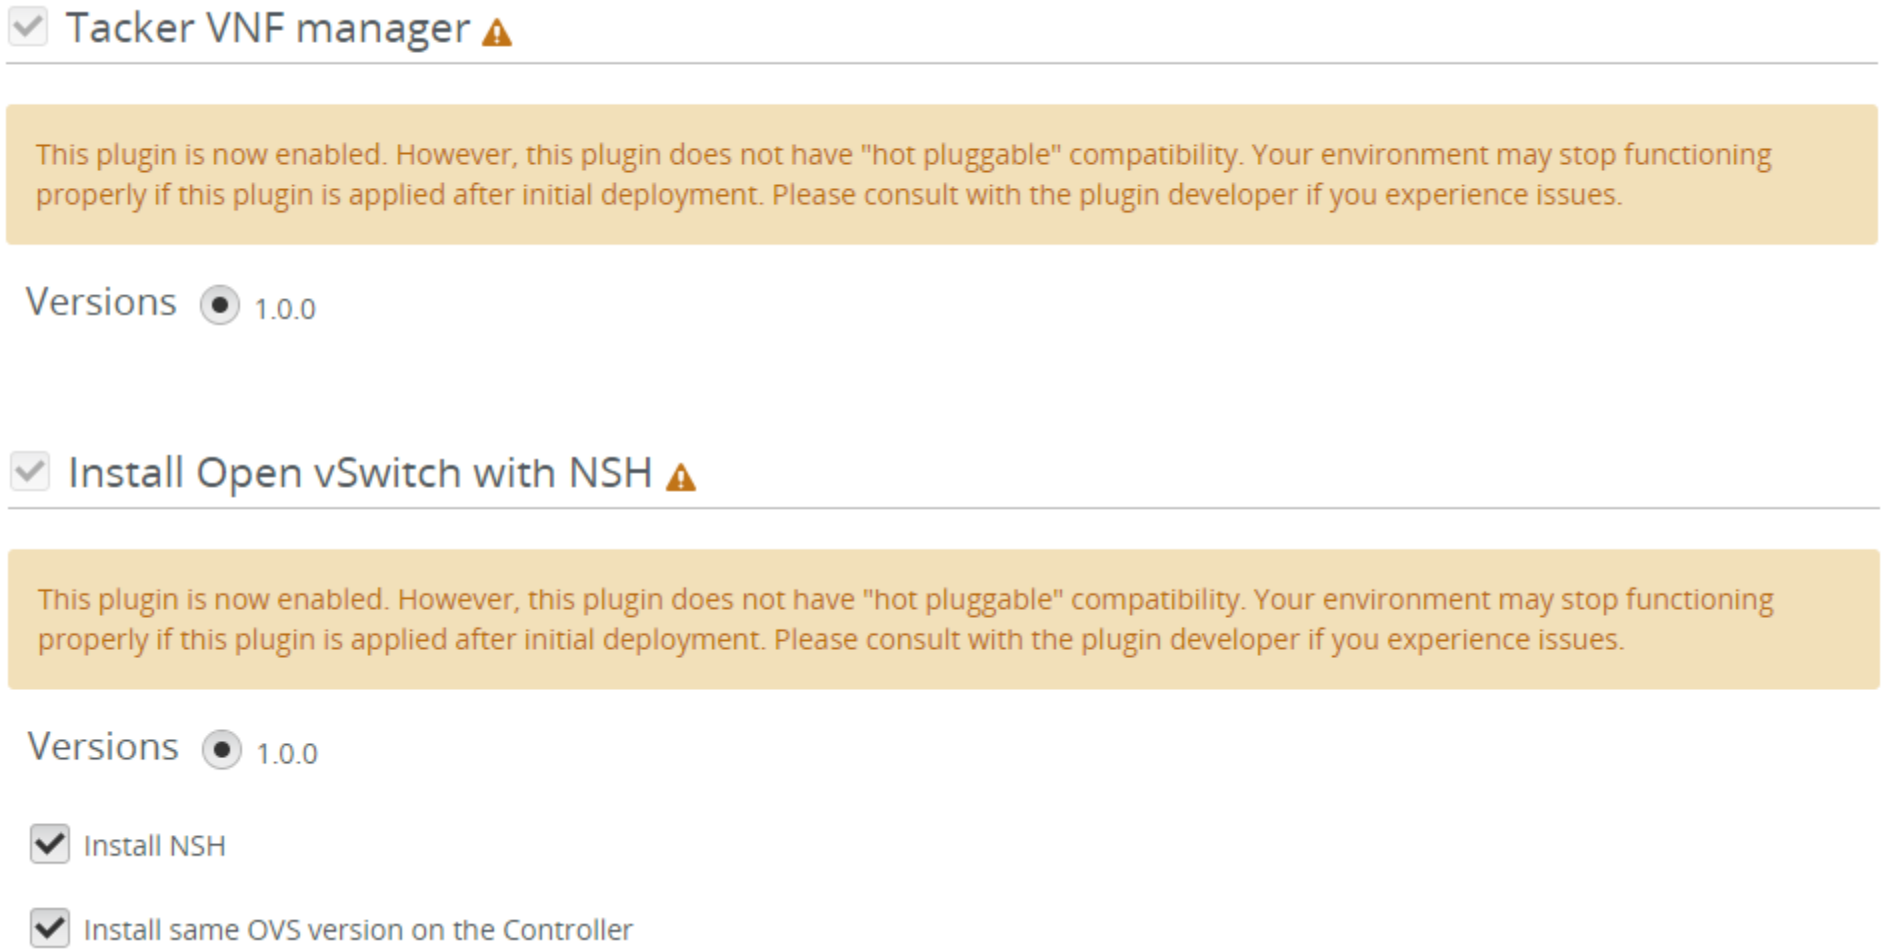

OPNFV的实践—Plugin

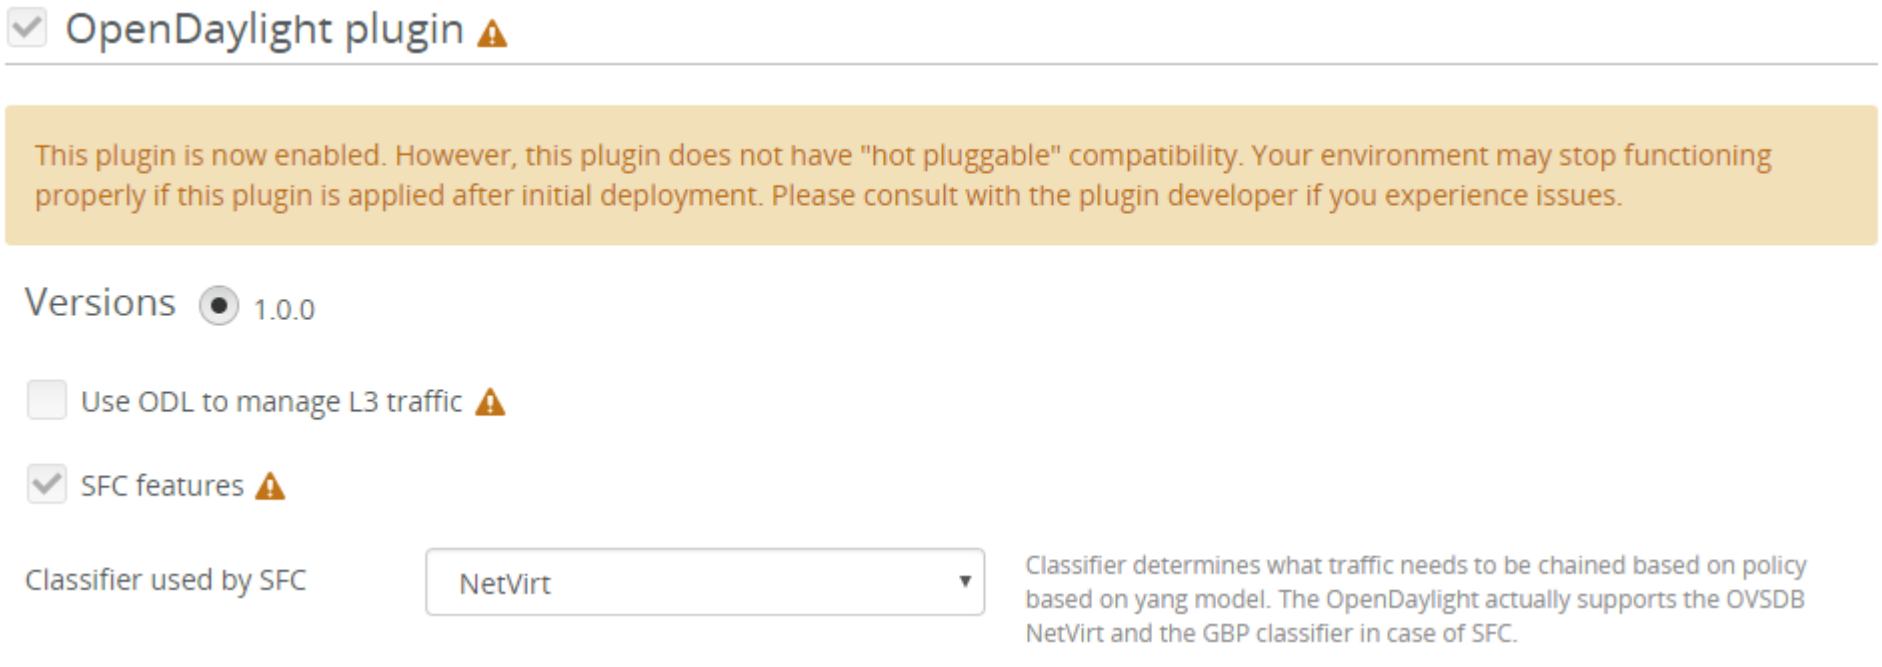

OPNFV的实践—ODL/SFC

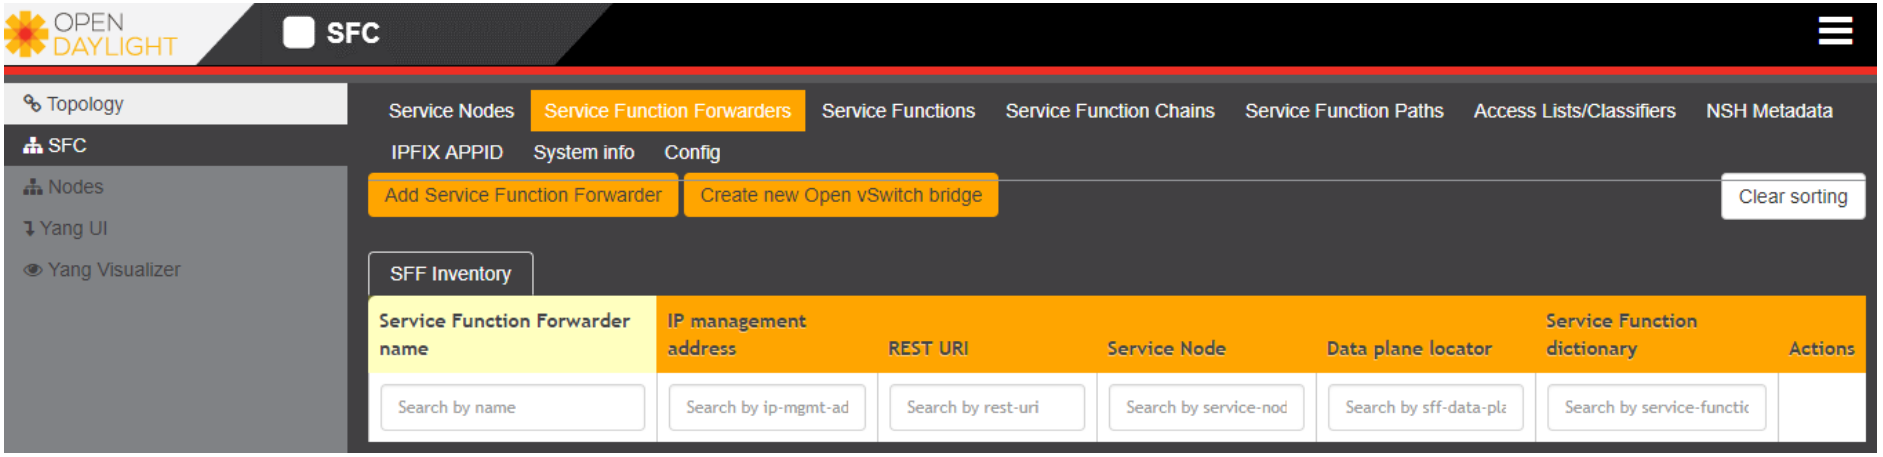

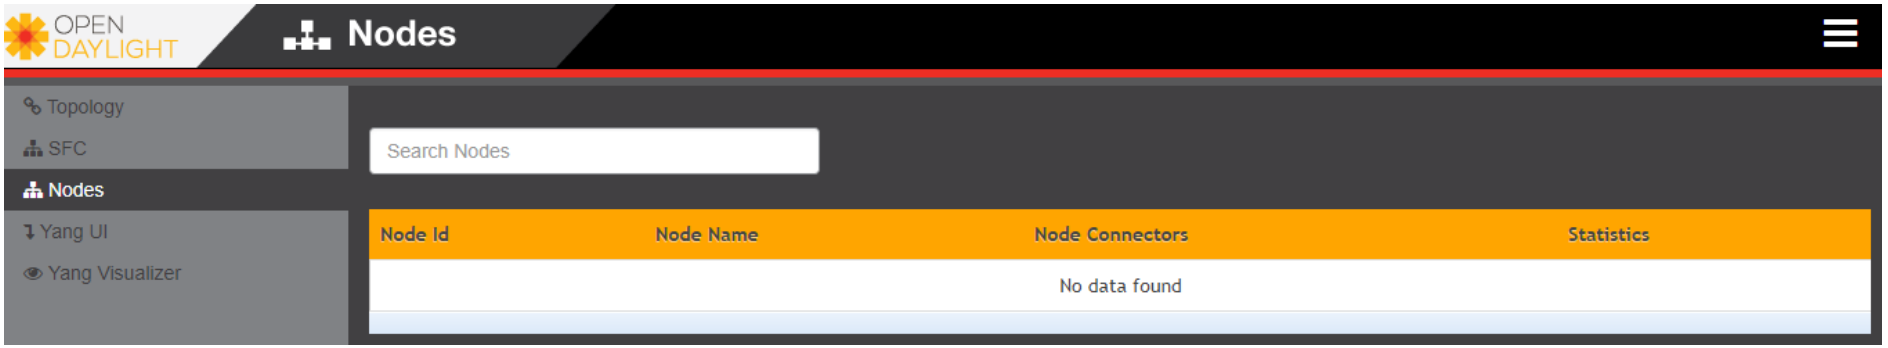

OPNFV的实践—NSH


```
root@node-2:~# ovs-vsctl list Open_vSwitch
_uuid               : 71312e3c-b7a9-4398-90ae-3c8768b57b03
bridges             : [79c1d4a0-0157-40f4-9064-b69b78d199a7]
cur_cfg             : 2
datapath_types      : [netdev, system]
db_version          : "7.14.0"
external_ids        : {hostname="node-2.domain.tld", system-id="c5772249-ecbc-4f8e-8705-0841ac71bd50"}
iface_types         : [geneve, gre, internal, ipsec_gre, lisp, patch, stt, system, tap, vxlan]
manager_options     : [805a0f7a-2ce1-42b1-b20e-cc34180874e1]
next_cfg            : 2
other_config        : {local_ip="192.168.4.1"}
ovs_version         : "2.6.1"
ssl                 : []
statistics          : {}
system_type         : ubuntu
system_version      : "16.04"
root@node-2:~#
```

OPNFV的实践—Tacker


```
Commands for API v1.0:
  bash-completion          Prints all of the commands and options for bash-completion.
  device-create            create a Device.
  device-delete            Delete a given Device.
  device-list              List device that belong to a given tenant.
  device-show              show information of a given Device.
  device-template-create   create a DeviceTemplate.
  device-template-delete   Delete a given DeviceTemplate.
  device-template-list     List device template that belong to a given tenant.
  device-template-show     show information of a given DeviceTemplate.
  device-template-update   Update a given DeviceTemplate.
  device-update            Update a given Device.
  ext-list                 List all extensions.
  ext-show                 Show information of a given resource.
  help                     print detailed help for another command
  interface-attach         Attach a network interface to a server.
  interface-detach         Detach a network interface from a server.
  sfc-classifier-create    Create a Service Function Chain Classifier
  sfc-classifier-delete    Delete a Service Function Chain
  sfc-classifier-list      List all Service Function Chain Classifiers
  sfc-classifier-show      Show a Service Function Chain
  sfc-classifier-update    Update a Service Function Chain
  sfc-create               Create a Service Function Chain
  sfc-delete               Delete a Service Function Chain
  sfc-list                 List all Service Function Chains
  sfc-show                 Show a Service Function Chain
  sfc-update               Update a Service Function Chain
  vnf-create               create a VNF.
  vnf-delete               Delete a given VNF.
  vnf-list                 List device that belong to a given tenant.
  vnf-show                 show information of a given VNF.
  vnf-update               Update a given VNF.
  vnfd-create              create a VNFD.
  vnfd-delete              Delete a given VNFD.
  vnfd-list                List VNFD that belong to a given tenant.
  vnfd-show                show information of a given VNFD.
root@node-1:~#
```

OPNFV的实践—Tacker


```
root@node-1:~# tacker vnfd-list
+--------------------------------------+------+-------------+--------------+-------------+
| id                                   | name | description | infra_driver | mgmt_driver |
+--------------------------------------+------+-------------+--------------+-------------+
| 90a79f81-76ae-40be-9803-ace87c20a91d |      | Demo example | heat         | noop        |
+--------------------------------------+------+-------------+--------------+-------------+
root@node-1:~# tacker vnf-list
+--------------------------------------+------+-------------+----------+--------+
| id                                   | name | description | mgmt_url | status |
+--------------------------------------+------+-------------+----------+--------+
| 29a2089f-279f-439b-9709-0cb7bd63e357 |      | Demo example |          | ACTIVE |
+--------------------------------------+------+-------------+----------+--------+
root@node-1:~# heat stack-list
WARNING (shell) "heat stack-list" is deprecated, please use "openstack stack list" instead
+--------------------------------------+-----------------------------------------------------------------------------------------------------+-----------------+----------------------+---------
-----+
| id                                   | stack_name                                                                                          | stack_status    | creation_time        | updated_
time |
+--------------------------------------+-----------------------------------------------------------------------------------------------------+-----------------+----------------------+---------
-----+
| 7801bc36-6418-4cf6-91e7-4017ab1d2436 | tacker.vm.drivers.heat.heat_DeviceHeat-29a2089f-279f-439b-9709-0cb7bd63e357 | CREATE_COMPLETE | 2017-05-21T07:37:09Z | None
     |
+--------------------------------------+-----------------------------------------------------------------------------------------------------+-----------------+----------------------+---------
-----+
root@node-1:~# nova list
+----+------+--------+------------+-------------+----------+
| ID | Name | Status | Task State | Power State | Networks |
+----+------+--------+------------+-------------+----------+
+----+------+--------+------------+-------------+----------+
root@node-1:~#
```

1 NFV架构介绍

2 OPNFV的实践

3 Tacker+Networking-SFC的实践

4 遗留问题

Tacker对Networking-SFC的驱动

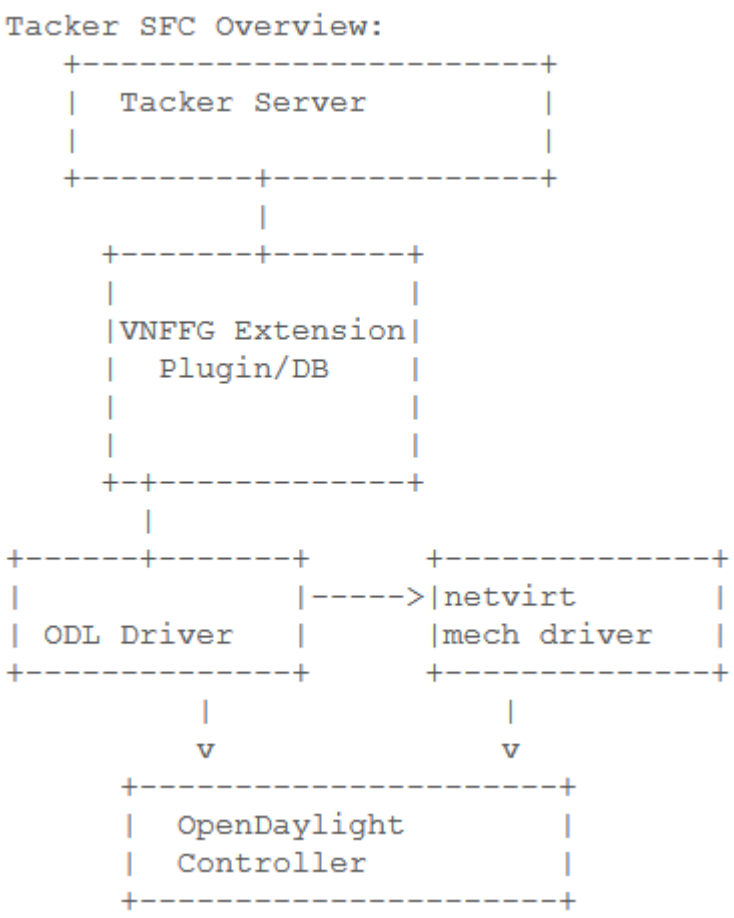

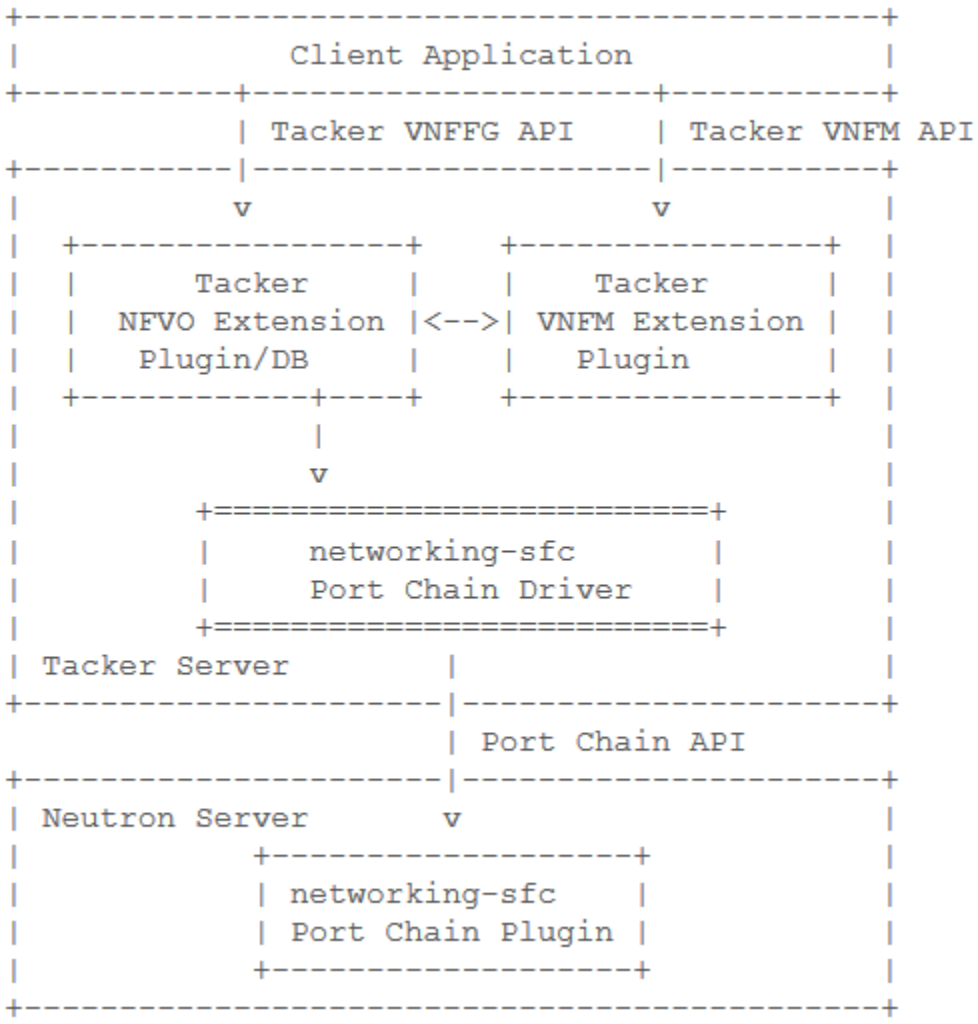

https://specs.openstack.org/openstack/tacker-specs/index.html

Tacker的安装


```
root@controller:/home/ubuntu# tacker vim-list
+----------------------+----------------------+------+----------+-------------+----------------------+----------------------+----------------------+-----------+
| id                   | tenant_id            | name | type     | description | auth_url             | placement_attr       | auth_cred            | status    |
+----------------------+----------------------+------+----------+-------------+----------------------+----------------------+----------------------+-----------+
| 0f6e6fc7-7684-41dc-  | a08b4c2905894b7891d07 | VIM  | openstack |             | http://10.90.3.55:353 | {u'regions':         | {u'username':        | REACHABLE |
| a5a6-aff7792905d9    | 1306f858f08          |      |          |             | 57/v3                | [u'RegionOne']}      | u'tacker', u'password': |           |
|                      |                      |      |          |             |                      |                      | u'***',              |           |
|                      |                      |      |          |             |                      |                      | u'project_name':     |           |
|                      |                      |      |          |             |                      |                      | u'service',          |           |
|                      |                      |      |          |             |                      |                      | u'user_domain_name': |           |
|                      |                      |      |          |             |                      |                      | u'default',          |           |
|                      |                      |      |          |             |                      |                      | u'auth_url': u'http://1 |           |
|                      |                      |      |          |             |                      |                      | 0.90.3.55:35357/v3', |           |
|                      |                      |      |          |             |                      |                      | u'project_id': None, |           |
|                      |                      |      |          |             |                      |                      | u'project_domain_name': |           |
|                      |                      |      |          |             |                      |                      | u'default'}          |           |
+----------------------+----------------------+------+----------+-------------+----------------------+----------------------+----------------------+-----------+
root@controller:/home/ubuntu# tacker vnfd-list
+--------------------------------------+-------------+---------------------+-------------+-------------+
| id                                   | name        | description         | infra_driver | mgmt_driver |
+--------------------------------------+-------------+---------------------+-------------+-------------+
| 1c19adae-cd98-4265-9b56-ce7db6588b1f | vfw-uplink  | vfw-uplink example  | heat         | noop        |
| 711c5d5e-27e3-4f8d-b8fb-9543f43b52cb | url-filter  | url-filter example  | heat         | noop        |
| 72e4b1ec-cce4-44fd-8e0b-6acb3b57cce8 | vfw-downlink | vfw-downlink example | heat         | noop        |
| f77f545b-cb0a-4069-b294-3bdb0baf684b | waf         | waf example         | heat         | noop        |
+--------------------------------------+-------------+---------------------+-------------+-------------+
root@controller:/home/ubuntu# tacker vnf-list
+----------------------+-------------+---------------------+----------------------+--------+----------------------+----------------------+--------------+
| id                   | name        | description         | mgmt_url             | status | vim_id               | placement_attr       | error_reason |
+----------------------+-------------+---------------------+----------------------+--------+----------------------+----------------------+--------------+
| 3a2aeb22-2117-46c3-8a4 | vfw-uplink | vfw-uplink example | {"VDU1": "10.90.3.91"} | ACTIVE | 0f6e6fc7-7684-41dc- | {u'vim_name': u'VIM'} |              |
| 2-6509f22fcea4       |             |                     |                      |        | a5a6-aff7792905d9    |                      |              |
| 553e7856-58fe-4377   | url-filter  | url-filter example  | {"VDU2":             | ACTIVE | 0f6e6fc7-7684-41dc- | {u'vim_name': u'VIM'} |              |
| -927d-c73d4608b70c   |             |                     | "172.16.11.4"}       |        | a5a6-aff7792905d9    |                      |              |
| 560e0ed4-8b66-45bc-a83 | waf        | waf example         | {"VDU1":             | ACTIVE | 0f6e6fc7-7684-41dc- | {u'vim_name': u'VIM'} |              |
| 8-f4900f67265b       |             |                     | "172.16.12.13"}      |        | a5a6-aff7792905d9    |                      |              |
| da7fcd76-06fc-       | vfw-downlink | vfw-downlink example | {"VDU1":             | ACTIVE | 0f6e6fc7-7684-41dc- | {u'vim_name': u'VIM'} |              |
| 43cd-8268-f8a478a987c2 |           |                     | "192.168.0.1"}       |        | a5a6-aff7792905d9    |                      |              |
+----------------------+-------------+---------------------+----------------------+--------+----------------------+----------------------+--------------+
```

https://docs.openstack.org/tacker/latest/install/manual_installation.html

Networking-SFC的安装

```
root@controller:/home/ubuntu# neutron flow-classifier-list
+--------------------------------------+------+-----------------------------------------------------------------+
| id                                   | name | summary                                                         |
+--------------------------------------+------+-----------------------------------------------------------------+
| 2c7bf2cc-b63a-4a3c-ac7f-e4d5b6e00b26 |      | protocol: TCP,                                                  |
|                                      |      | source[port]: any[any:any],                                     |
|                                      |      | destination[port]: 0.0.0.0/0[80:80],                            |
|                                      |      | neutron_source_port: 20a1e096-8e08-4ddd-ac13-3e44e3d4deae,      |
|                                      |      | neutron_destination_port: None,                                 |
|                                      |      | l7_parameters: {}                                               |
| bb39beb5-4f26-4ccc-a71d-2b7fed6d1d9b |      | protocol: TCP,                                                  |
|                                      |      | source[port]: any[any:any],                                     |
|                                      |      | destination[port]: 192.168.0.250/32[8080:8080],                 |
|                                      |      | neutron_source_port: 97815697-f23e-4261-abc9-ff06a402e187,      |
|                                      |      | neutron_destination_port: None,                                 |
|                                      |      | l7_parameters: {}                                               |
+--------------------------------------+------+-----------------------------------------------------------------+
root@controller:/home/ubuntu# neutron port-pair-list
+--------------------------------------+-----------------------------+--------------------------------------+--------------------------------------+
| id                                   | name                        | ingress                              | egress                               |
+--------------------------------------+-----------------------------+--------------------------------------+--------------------------------------+
| 4153229a-4691-49e3-a550-4679e992eb9f | waf-connection-points       | 1573f9c6-8a7e-47c2-9239-c40909c4715c | 1573f9c6-8a7e-47c2-9239-c40909c4715c |
| 87dfb56c-3bc5-404b-b781-fe5151a534a1 | url-filter-connection-points | 8d8c0e72-0fd5-4c00-8ce2-97e3ee3fc1ee | 8d8c0e72-0fd5-4c00-8ce2-97e3ee3fc1ee |
+--------------------------------------+-----------------------------+--------------------------------------+--------------------------------------+
root@controller:/home/ubuntu# neutron port-pair-group-list
+--------------------------------------+---------------------------+----------------------------------------+
| id                                   | name                      | port_pairs                             |
+--------------------------------------+---------------------------+----------------------------------------+
| cea00bc6-6be9-4e18-b0b7-532552a9459c | waf-port-pair-group       | [u'4153229a-4691-49e3-a550-4679e992eb9f'] |
| ede6757c-61c5-47a9-8d6a-f8e411ca8d0f | url-filter-port-pair-group | [u'87dfb56c-3bc5-404b-b781-fe5151a534a1'] |
+--------------------------------------+---------------------------+----------------------------------------+
root@controller:/home/ubuntu# neutron port-chain-list
+--------------------------------------+-------------------------+----------------------------------------+----------------------------------------+
| id                                   | name                    | port_pair_groups                       | flow_classifiers                       |
+--------------------------------------+-------------------------+----------------------------------------+----------------------------------------+
| 66a7ea9a-fc84-4423-bcee-e55adfdac157 | sfc-downlink-port-chain | [u'cea00bc6-6be9-4e18-b0b7-532552a9459c'] | [u'bb39beb5-4f26-4ccc-a71d-2b7fed6d1d9b'] |
| ffe59229-e18b-47f8-a286-1ee99490a927 | sfc-uplink-port-chain   | [u'ede6757c-61c5-47a9-8d6a-f8e411ca8d0f'] | [u'2c7bf2cc-b63a-4a3c-ac7f-e4d5b6e00b26'] |
+--------------------------------------+-------------------------+----------------------------------------+----------------------------------------+
```

https://docs.openstack.org/networking-sfc/latest/installation.html

Networking-SFC的流表


```
root@controller:/home/ubuntu# tacker vnffgd-list
+--------------------------------------+--------------+-------------+
| id                                   | name         | description |
+--------------------------------------+--------------+-------------+
| 9c99475a-5d4c-42ca-b706-d83e426bffea | sfc-uplink   |             |
| f6d844c4-0429-48c1-b3d9-1f1e879ea283 | sfc-downlink |             |
+--------------------------------------+--------------+-------------+
root@controller:/home/ubuntu# tacker vnffg-list
+--------------------------------------+--------------+-------------+--------+--------------------------------------+
| id                                   | name         | description | status | vnffgd_id                            |
+--------------------------------------+--------------+-------------+--------+--------------------------------------+
| 6ab2c261-fa91-4712-9f6c-8f4d25bc8c20 | sfc-downlink |             | ACTIVE | f6d844c4-0429-48c1-b3d9-1f1e879ea283 |
| 7430142c-f0dd-4714-91fe-bb5001689fca | sfc-uplink   |             | ACTIVE | 9c99475a-5d4c-42ca-b706-d83e426bffea |
+--------------------------------------+--------------+-------------+--------+--------------------------------------+
root@controller:/home/ubuntu# ovs-ofctl dump-flows br-int | grep group
 cookie=0xbbdb70444a02889f, duration=1774.995s, table=0, n_packets=0, n_bytes=0, idle_age=1774, priority=30,tcp,in_port=13,nw_dst=192.168.0.250,tp_dst=8080 actions=grou
p:1
 cookie=0xbbdb70444a02889f, duration=1756.347s, table=0, n_packets=0, n_bytes=0, idle_age=1756, priority=30,tcp,in_port=18,tp_dst=80 actions=group:2
root@controller:/home/ubuntu# ovs-ofctl dump-groups br-int -O OpenFlow13
OFPST_GROUP_DESC reply (OF1.3) (xid=0x2):
 group_id=1,type=select,bucket=actions=set_field:fa:16:3e:3c:15:a2->eth_dst,resubmit(,5)
 group_id=2,type=select,bucket=actions=set_field:fa:16:3e:ea:c0:2b->eth_dst,resubmit(,5)
root@controller:/home/ubuntu#
```

NFV一体机服务链方案PoC

TOSCA （Topology Orchestration Specification for Cloud Applications）

管理员

VNF模板

OpenStack

Tacker+Networking-SFC

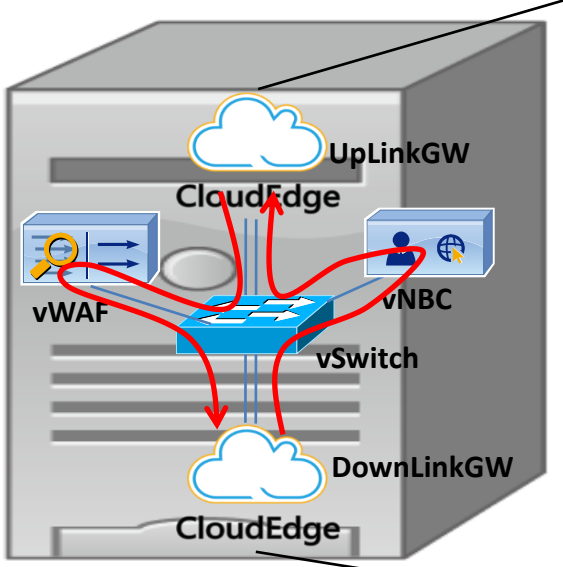

外部网

内部网

```
topology_template:
  description: Sample VNFFG template

  node_templates:

    Forwarding_path1:
      type: tosca.nodes.nfv.FP.Tacker
      description: creates path (ext->waf->internal)
      properties:
        id: 51
        policy:
          type: ACL
          criteria:
            - network_src_port_id: c299ebbb-1993-4019-b06b-ca024d7e6d01
            - ip_dst_prefix: 192.168.0.0/24
            - destination_port_range: 80-80
            - ip_proto: 6
        path:
          - forwarder: waf
            capability: CP12

  groups:
    VNFFG1:
      type: tosca.groups.nfv.VNFFG
      description: ext to internal
      properties:
        vendor: tacker
        version: 1.0
        number_of_endpoints: 5
        dependent_virtual_link: [VL12]
        connection_point: [CP12]
        constituent_vnfs: [waf]
      members: [Forwarding_path1]
```

流量入口

引流规则

服务节点

NFV业务操作—注册VIM

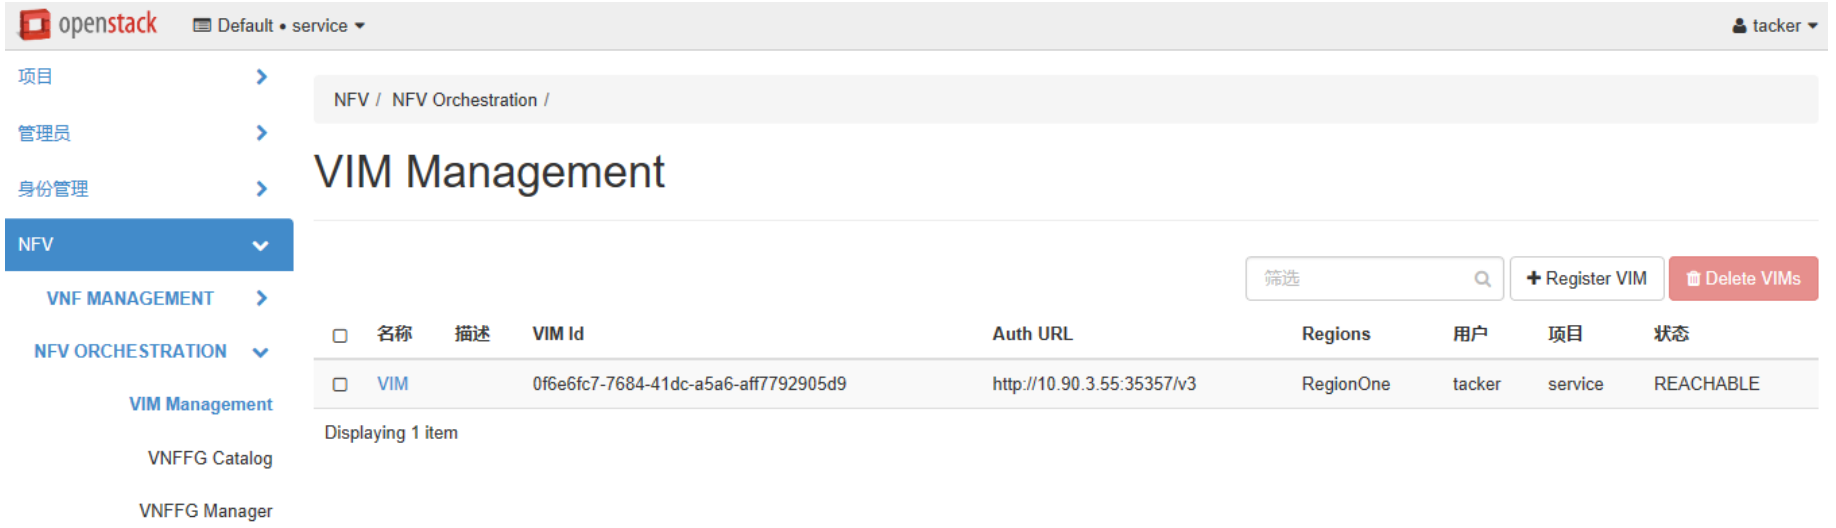

NFV业务操作—创建VNF模板

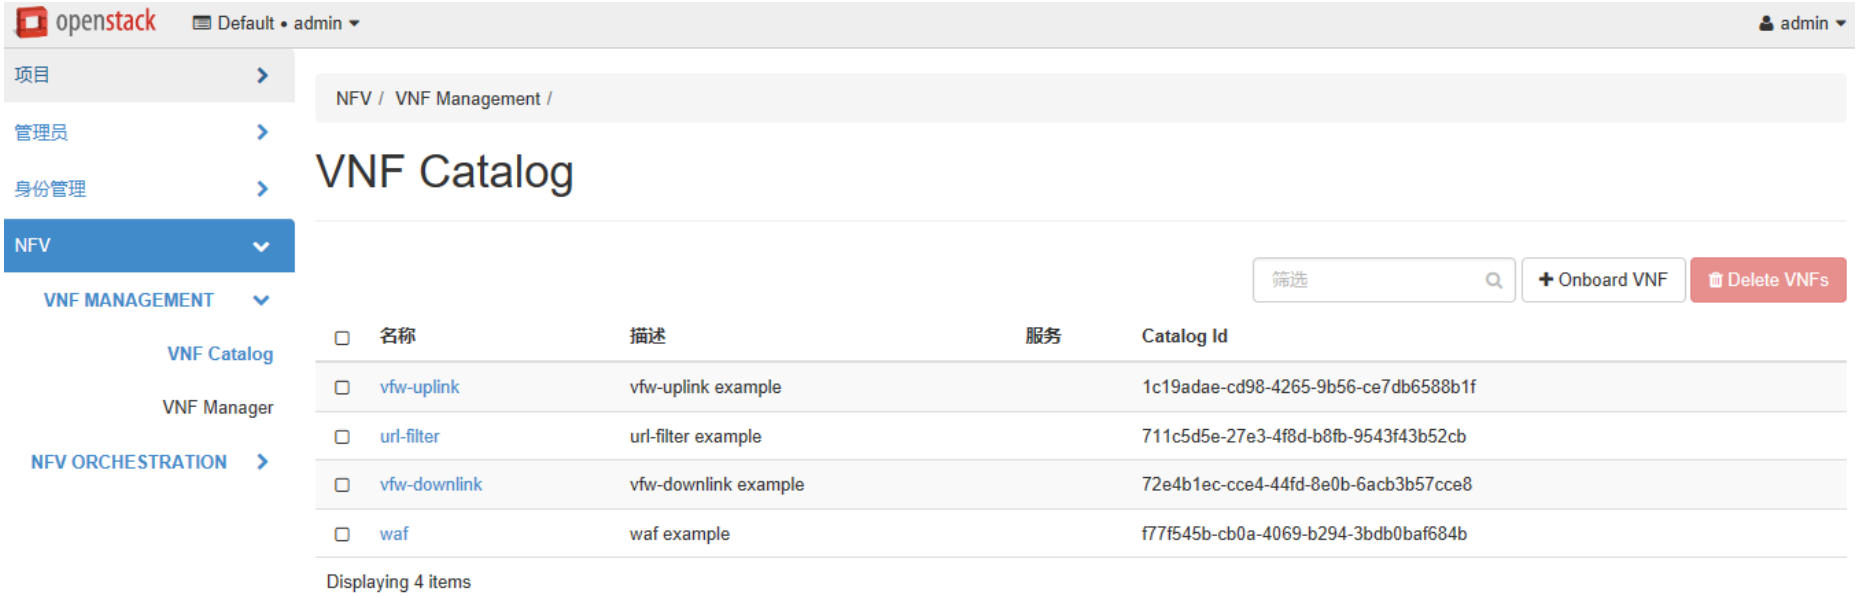

https://github.com/openstack/tacker/tree/master/samples/tosca-templates/vnfd

NFV业务操作—创建VNF实例

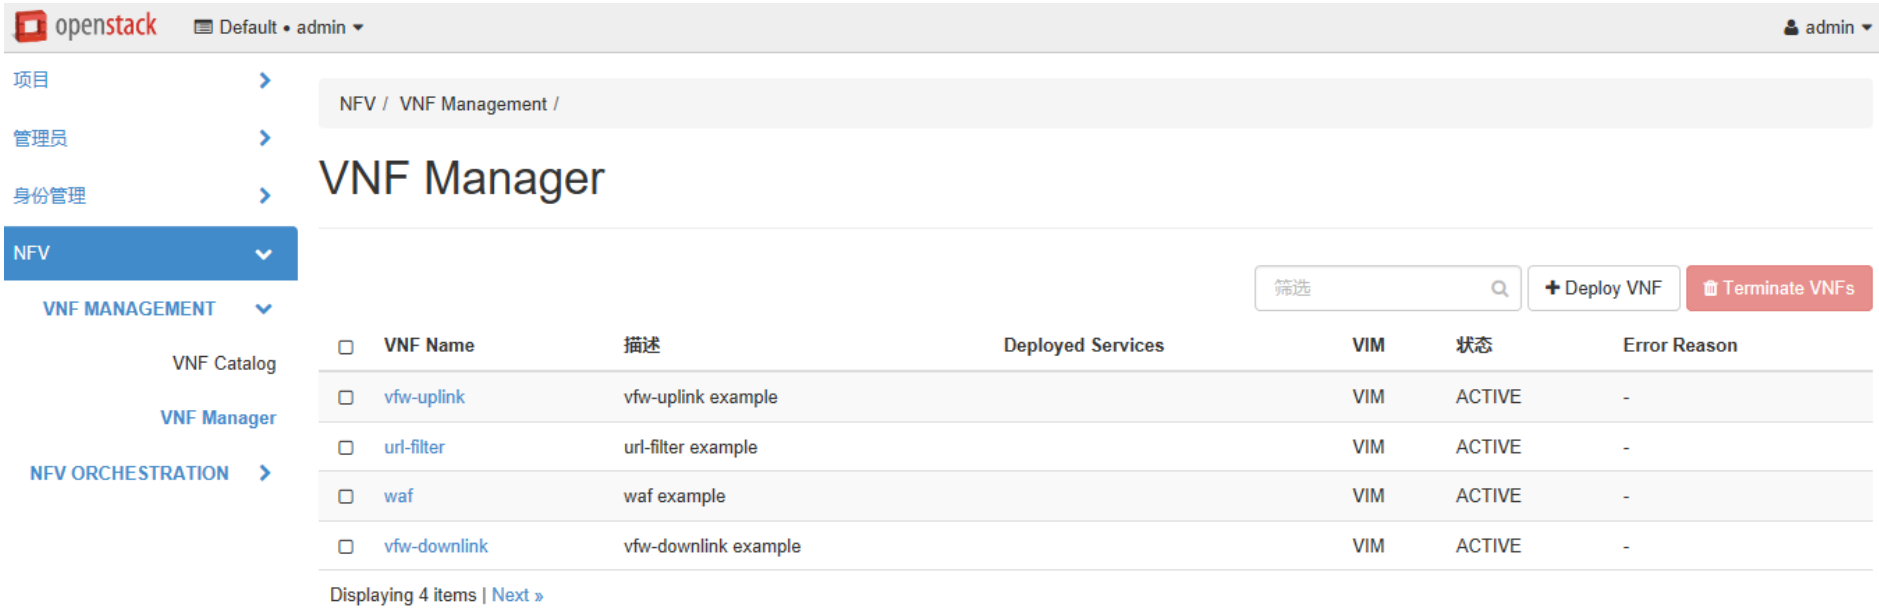

NFV业务操作—创建VNFFG模板

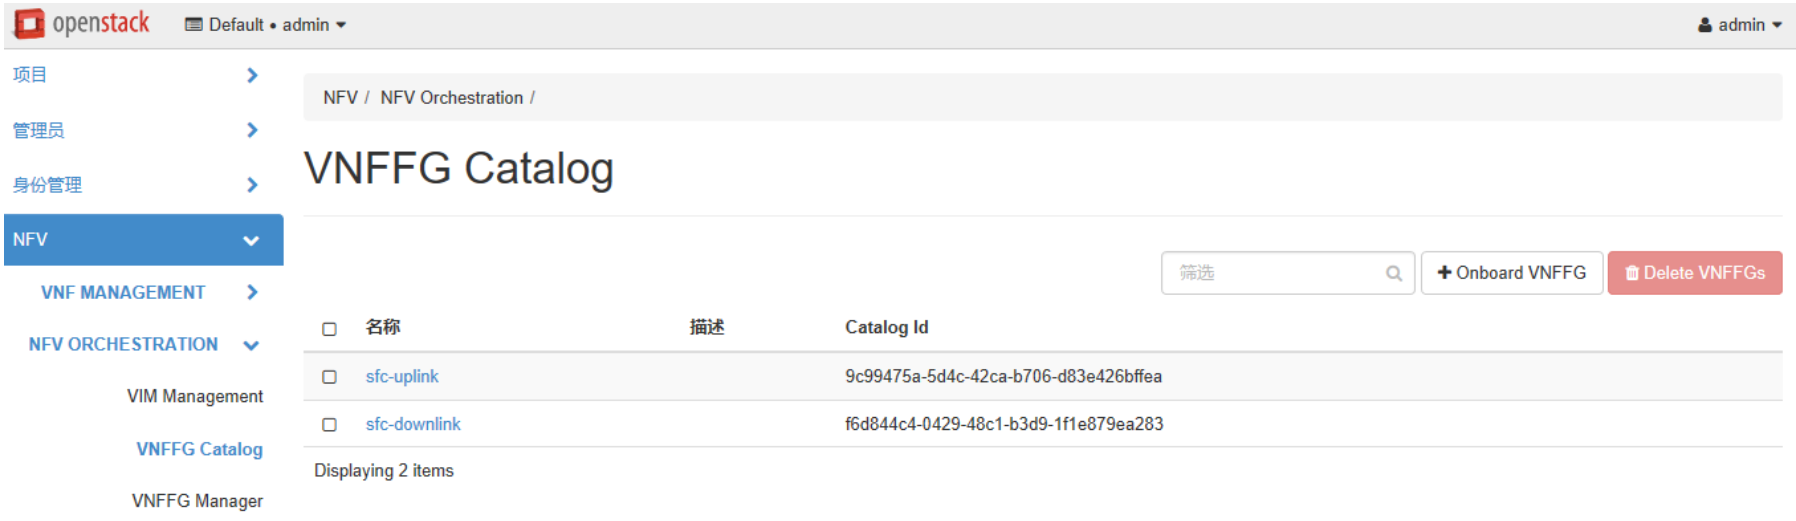

https://github.com/openstack/tacker/tree/master/samples/tosca-templates/vnffgd

NFV业务操作—创建VNFFG实例

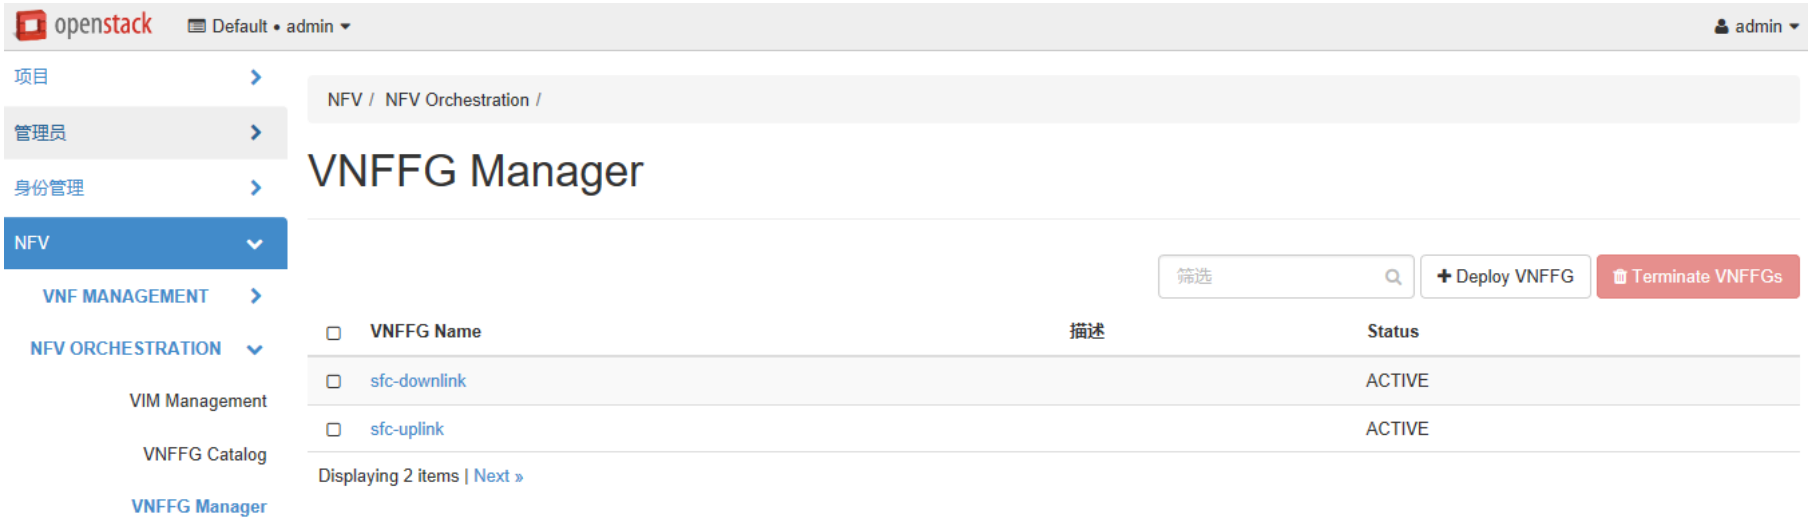

NFV业务操作—查看VNF栈

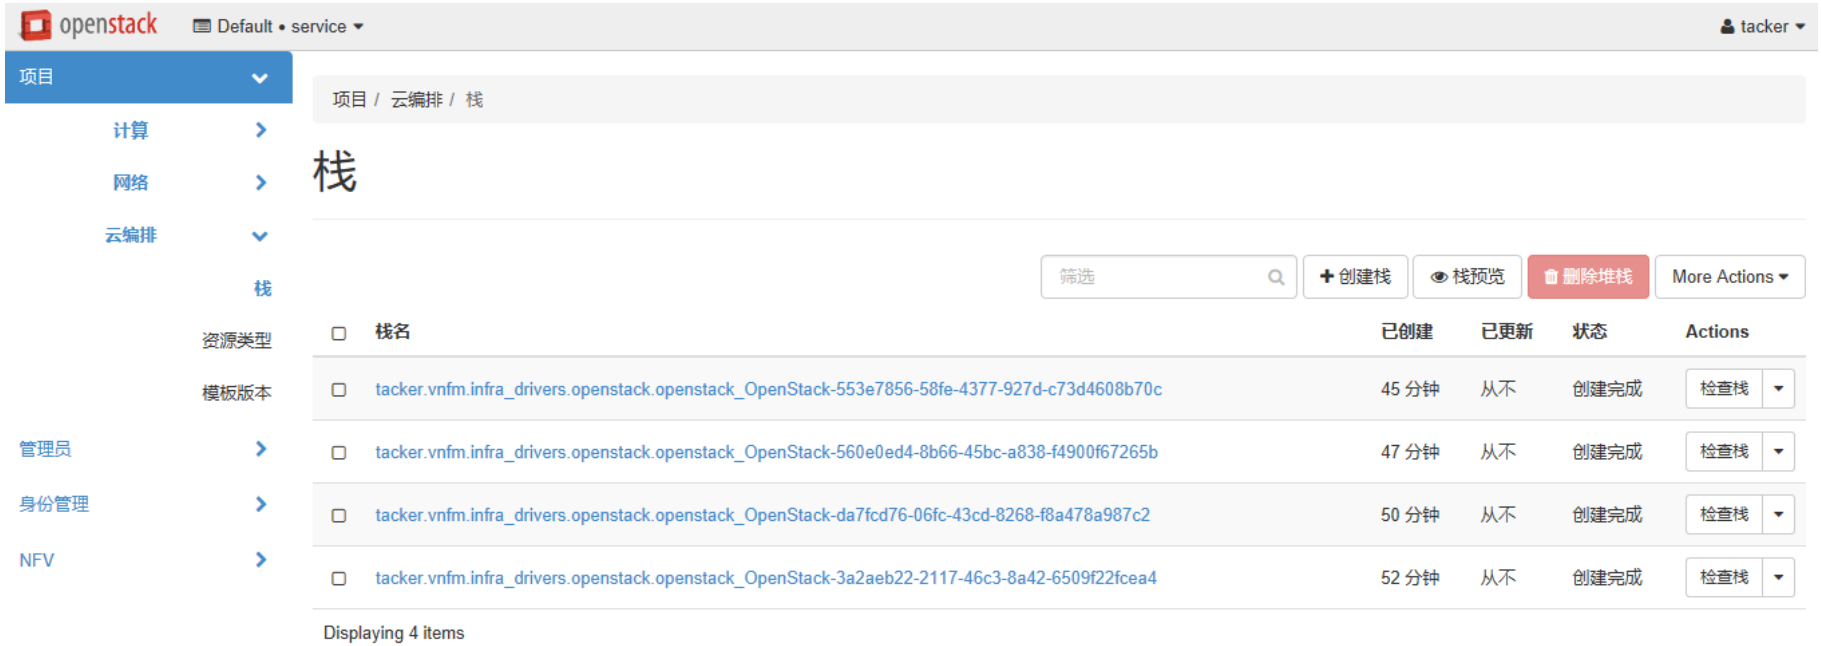

NFV业务操作—查看VNF虚机

1 NFV架构介绍

2 OPNFV的实践

3 Tacker+Networking-SFC的实践

4 遗留问题

遗留问题

- Tacker的Newton Stable版本有严重bug，界面上创建VNFFG失败，可以用Ocata Stable版本的vnffgmanager进行替换，仍然存在bug，界面上创建VNFFG时提示找不到VNF，但能成功。
- 普通Linux发行版不适合做VNF，不支持单臂路由。
- Ubuntu第二块网卡不抢占默认路由，Centos第二块网卡抢占默认路由。
- Ocata版本的Tacker+Networking-SFC虽然更成熟，但不支持嵌套虚拟化（libvert从1.3版本升级到2.5版本），给方案验证带来麻烦。

THANK YOU