



NFV PoC #43

Towards an efficient Data Plane processing

Jokin Garay <jokin.garay@keynetic.es>

Keynetic Technologies

2017/03/21

PoC #43: Team Members



- Network Operator / Service Provider
 - OSM & OpenVIM:
- Manufacturer – NFVI provider:



- Manufacturer – VNF FlowNAC:



- Additional members (research):



Keynetic Technologies

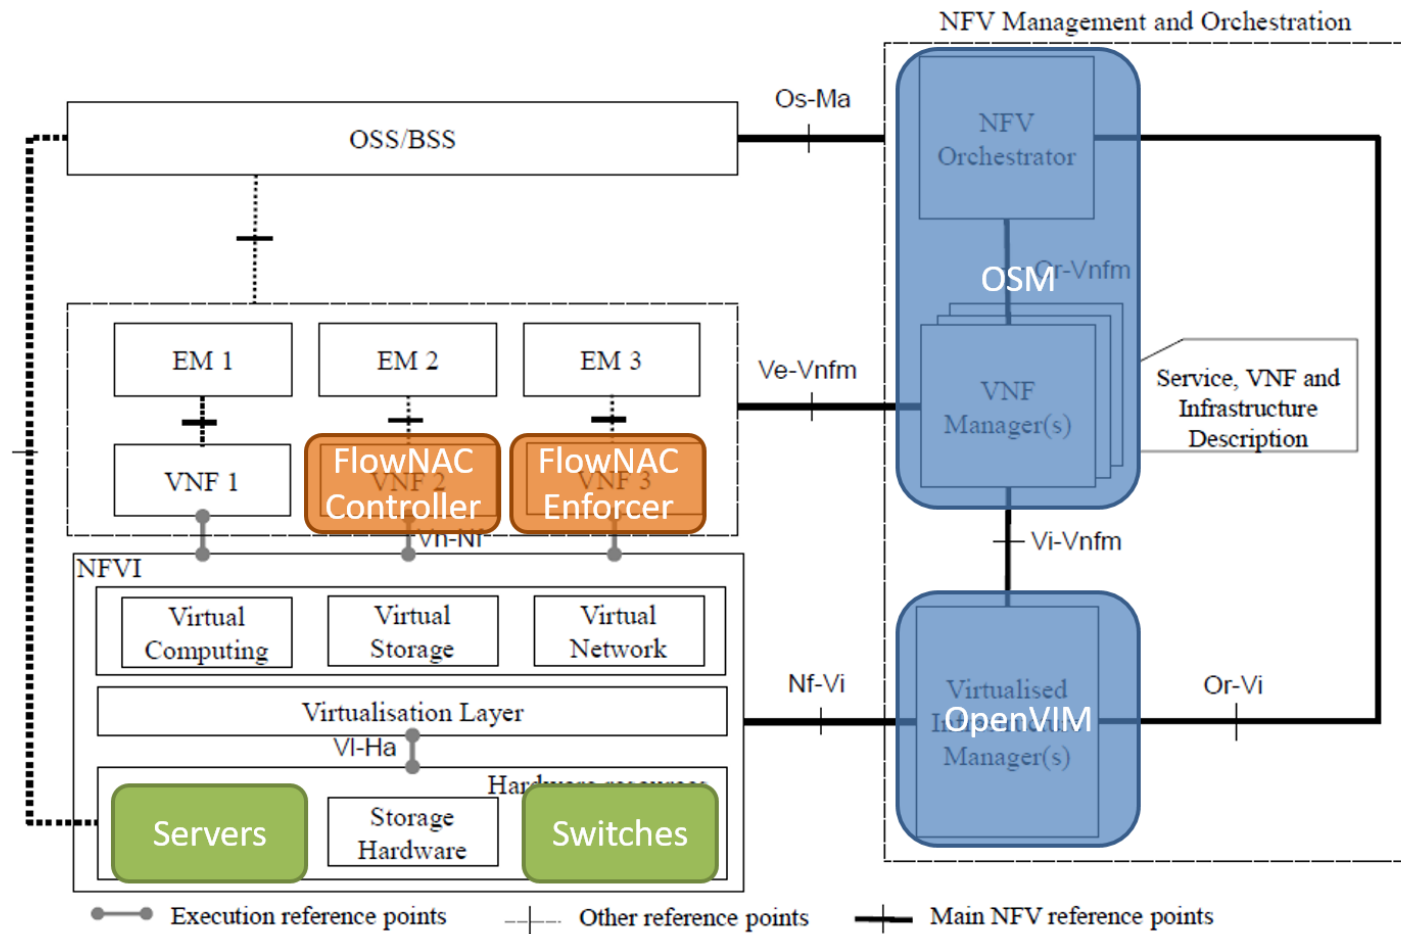


- Recently incorporated spinoff from the UPV/EHU I2T Research Group
- Background in R&D at european level and System Integration & Technology consulting
- Cybersecurity and network virtualization solutions based on SDN and NFV
- Involved in OSM since the kick off...
... and we've been in the Plugtest too

PoC #43: Mapping to NFV ISG Work

Component provided by


- Telefónica
- Keynetic - UPV/EHU
- HPE



http://nfvwiki.etsi.org/index.php?title=Toward_an_efficient_dataplane_processing

PoC #43: Public demonstration

- NFV#17 meeting (22nd - 25th February 2017)

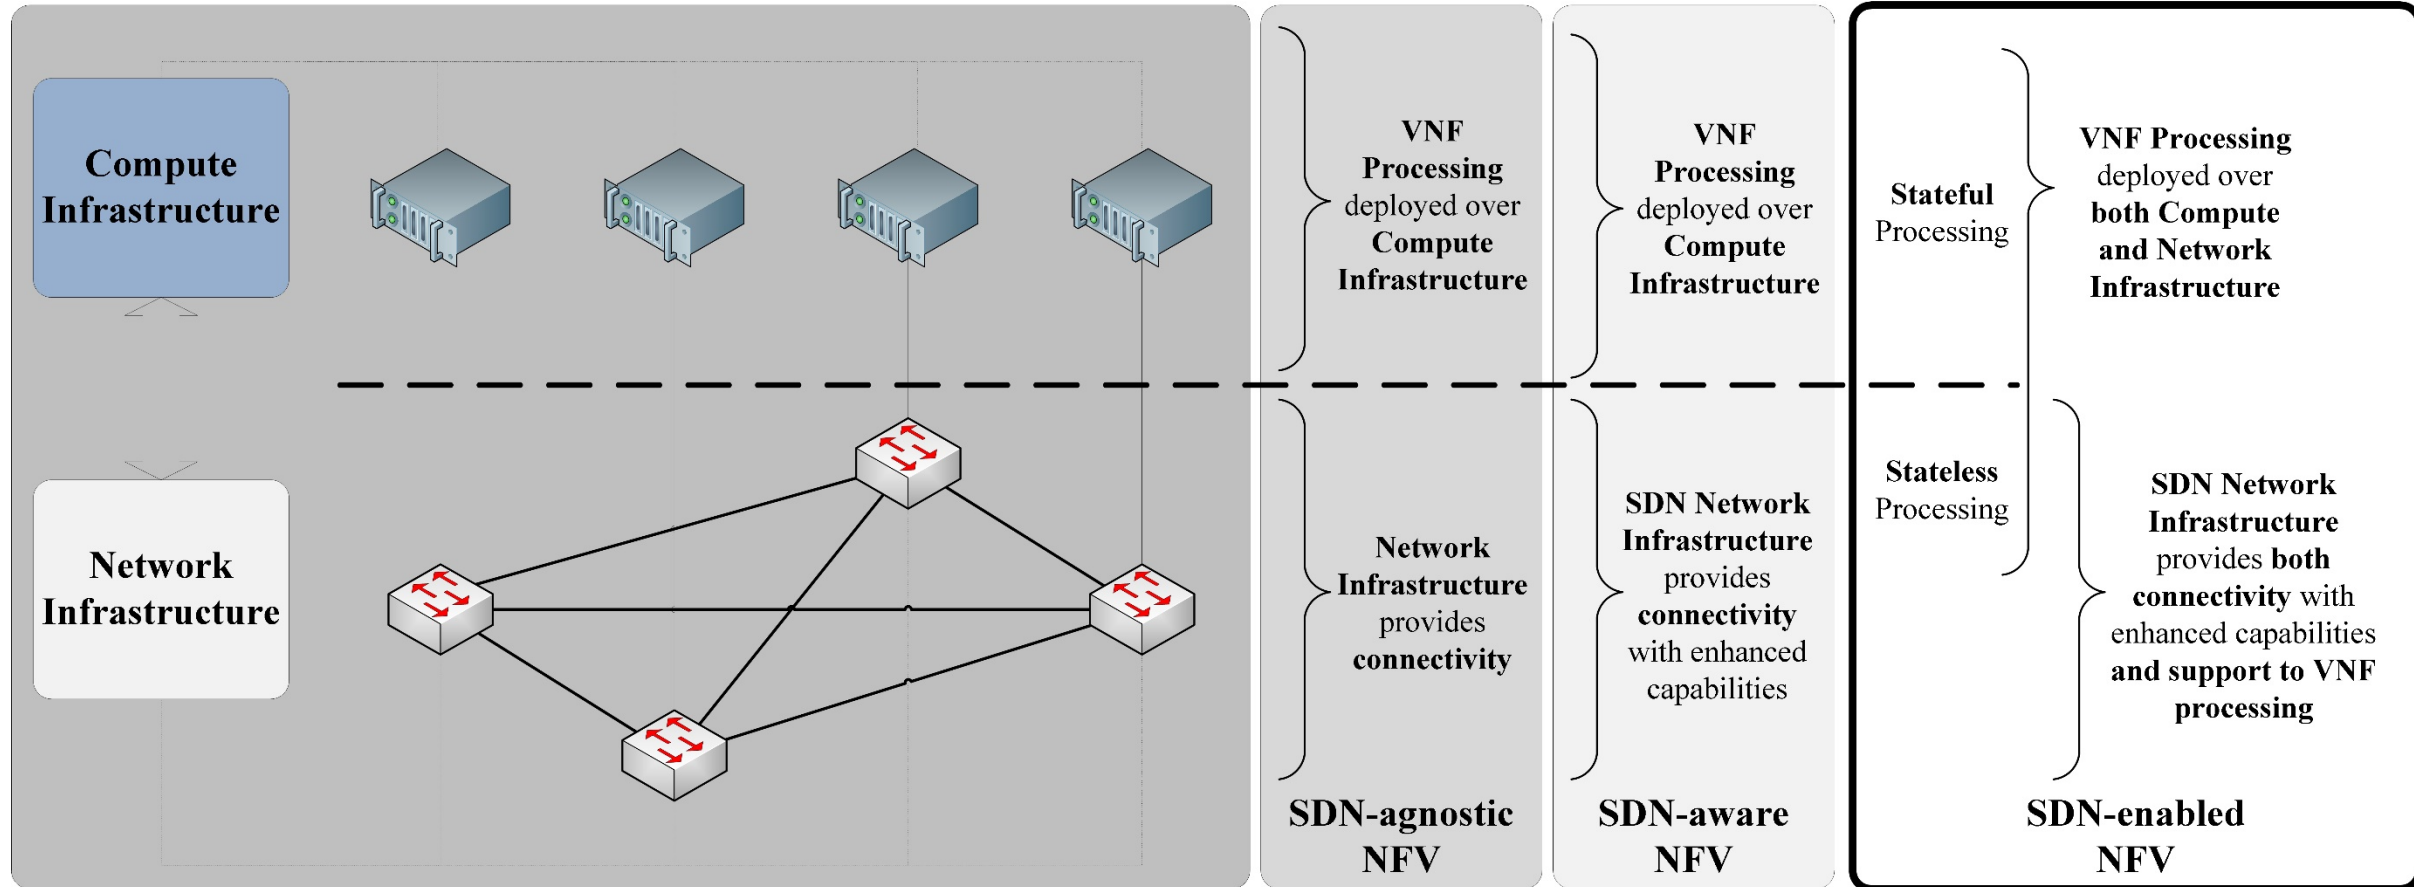


PoC #43: Goals



- VNF designed with separate sub-components:
 - Stateful network function components (compute)
... and stateless data path processing components (networking)
 - Allow each sub-component to independently scale
- Improvement of VNF data processing efficiency while minimizing the overall NFVI resources used
- Evaluate different alternatives to deploy the stateless data path processing component
 - Propose pointers to ease the integration into the ETSI NFV architectural framework

Towards SDN-enabled NFV

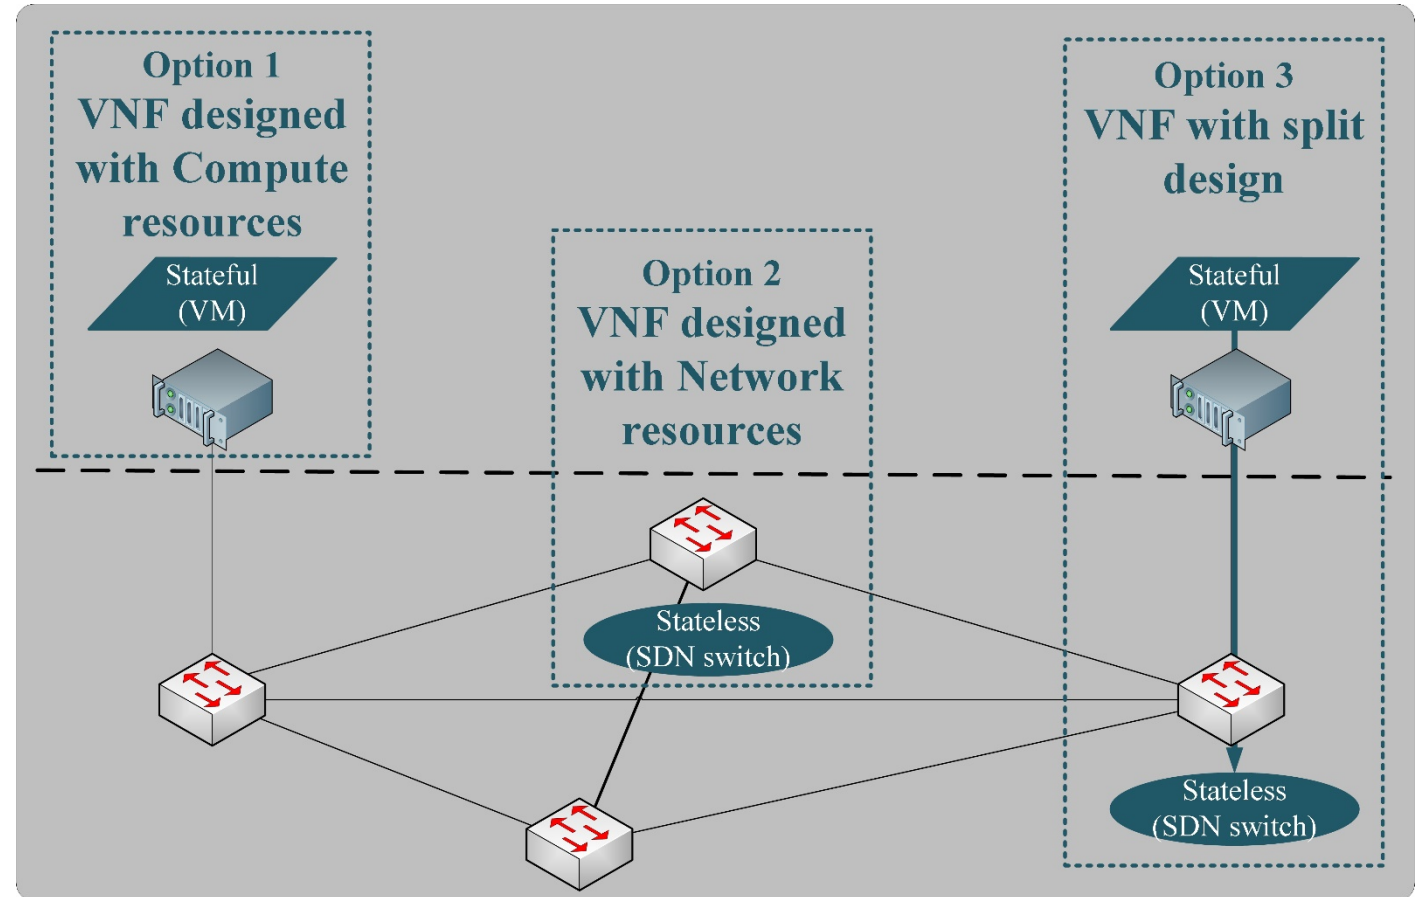



Matias, J., Garay, J., Toledo, N., Unzilla, J. and Jacob, E., 2015.

Toward an SDN-enabled NFV architecture. *IEEE Communications Magazine*, 53(4), pp.187-193.

<https://doi.org/10.1109/MCOM.2015.7081093>

Towards SDN-enabled NFV



Matias, J., Garay, J., Toledo, N., Unzilla, J. and Jacob, E., 2015.

Toward an SDN-enabled NFV architecture. *IEEE Communications Magazine*, 53(4), pp.187-193.

<https://doi.org/10.1109/MCOM.2015.7081093>

Towards SDN-enabled NFV



- Challenges arising from
 - VNF must be designed splitting the components to be deployed over compute and network resources
 - the network infrastructure must support a dual role for traffic steering and VNF processing
- Services must be redesigned, determine if the performance gain overcomes the effort involved
- Adds complexity to the optimal placement decision. The NFV framework must now orchestrate an additional type of resource with its own constraints
- **Bring Network Virtualization to the same level as Compute Virtualization**
 - Dual use of the infrastructure, for traffic steering and stateless NF processing, requires the underlying network infrastructure to guarantee isolation between both functionalities
 - Performance isolation is also required to avoid VNF functionality hindering the correct behavior of the NFV architecture

Matias, J., Garay, J., Toledo, N., Unzilla, J. and Jacob, E., 2015.


<https://doi.org/10.1109/MCOM.2015.7081093>

Toward an SDN-enabled NFV architecture. *IEEE Communications Magazine*, 53(4), pp.187-193.


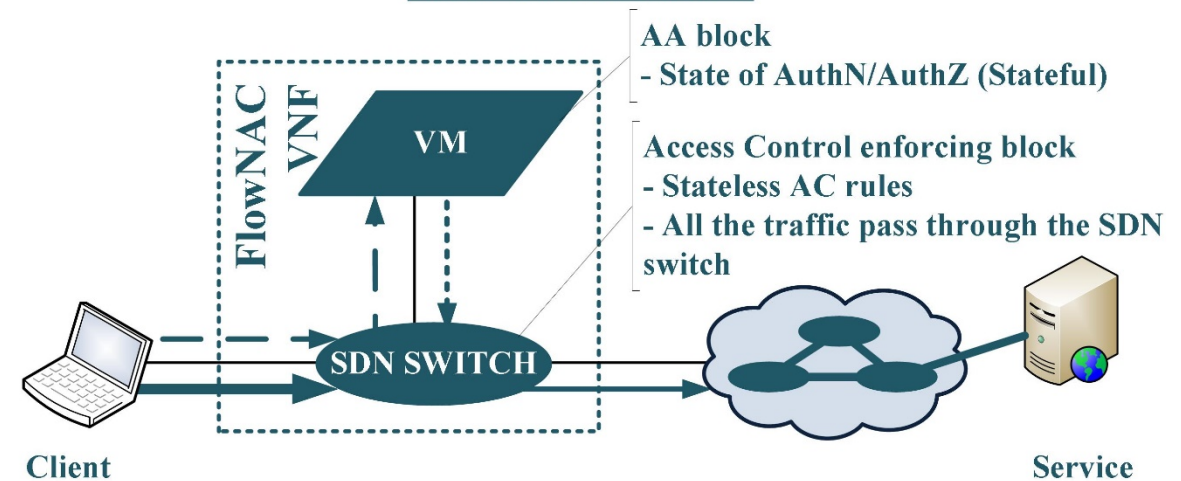
PoC #43: Technical Details

- Flow-based Network Access Control VNF

SDN-agnostic NFV / SDN-aware NFV



SDN-enabled NFV



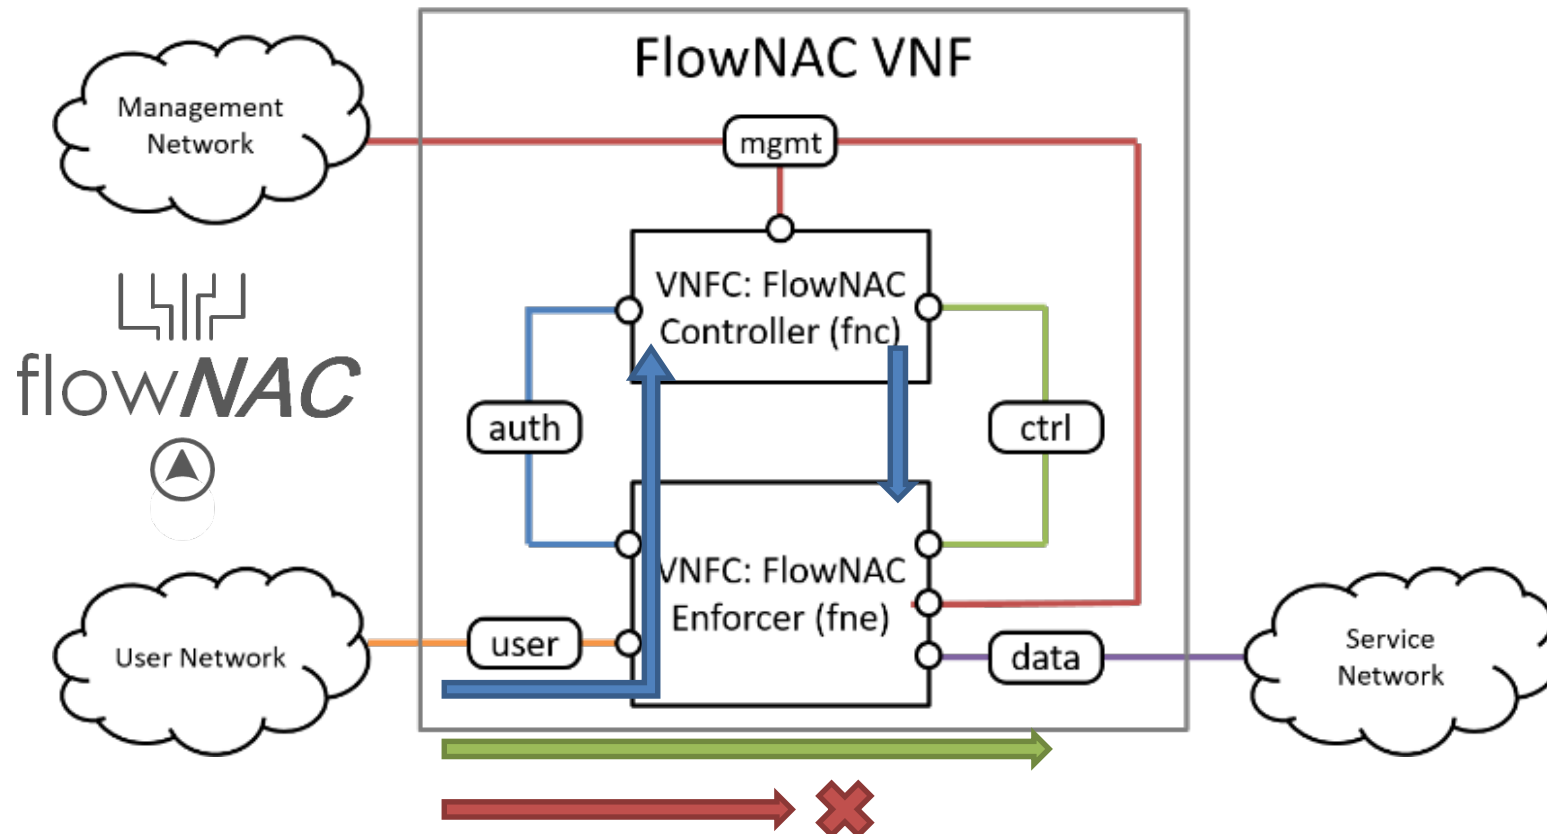
Matias, J., Garay, J., Toledo, N., Unzilla, J. and Jacob, E., 2015.

Toward an SDN-enabled NFV architecture. *IEEE Communications Magazine*, 53(4), pp.187-193.

<https://doi.org/10.1109/MCOM.2015.7081093>

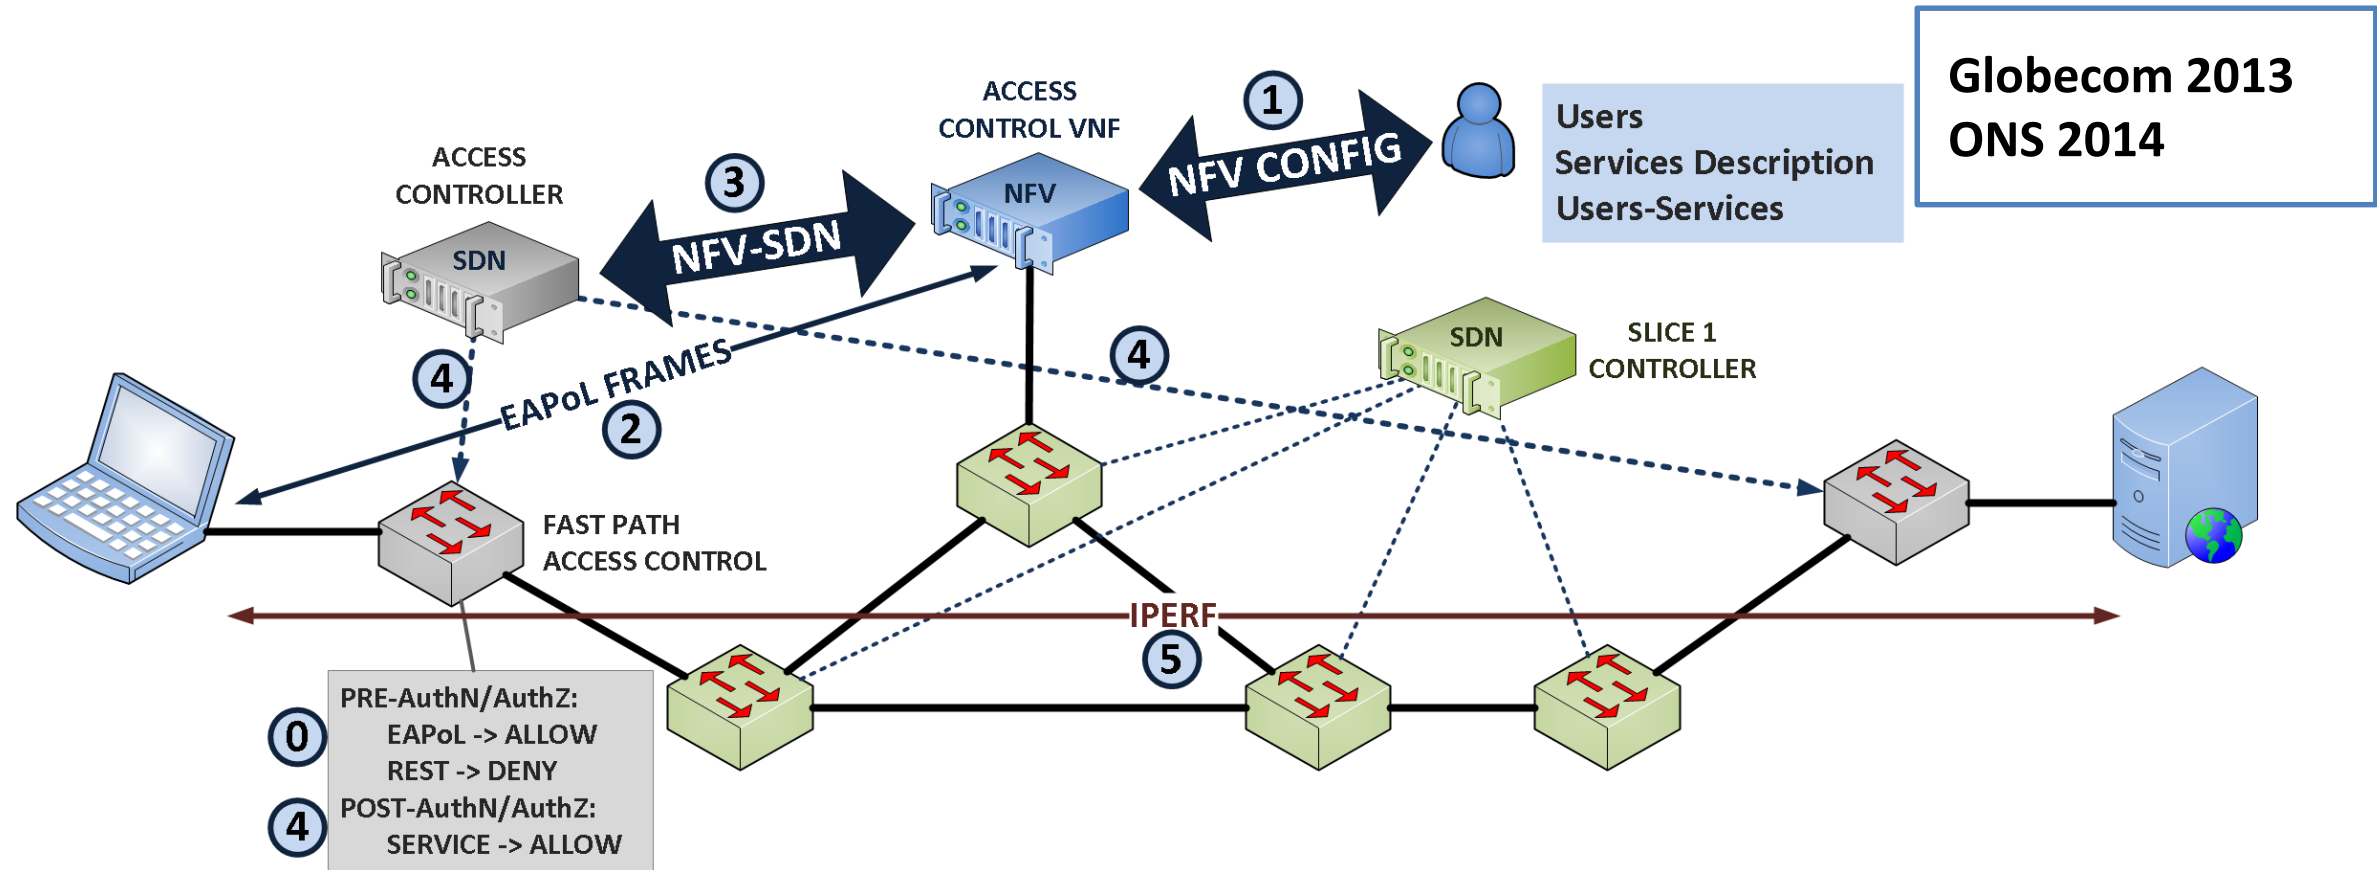
PoC #43: Technical Details

- Flow-based Network Access Control VNF



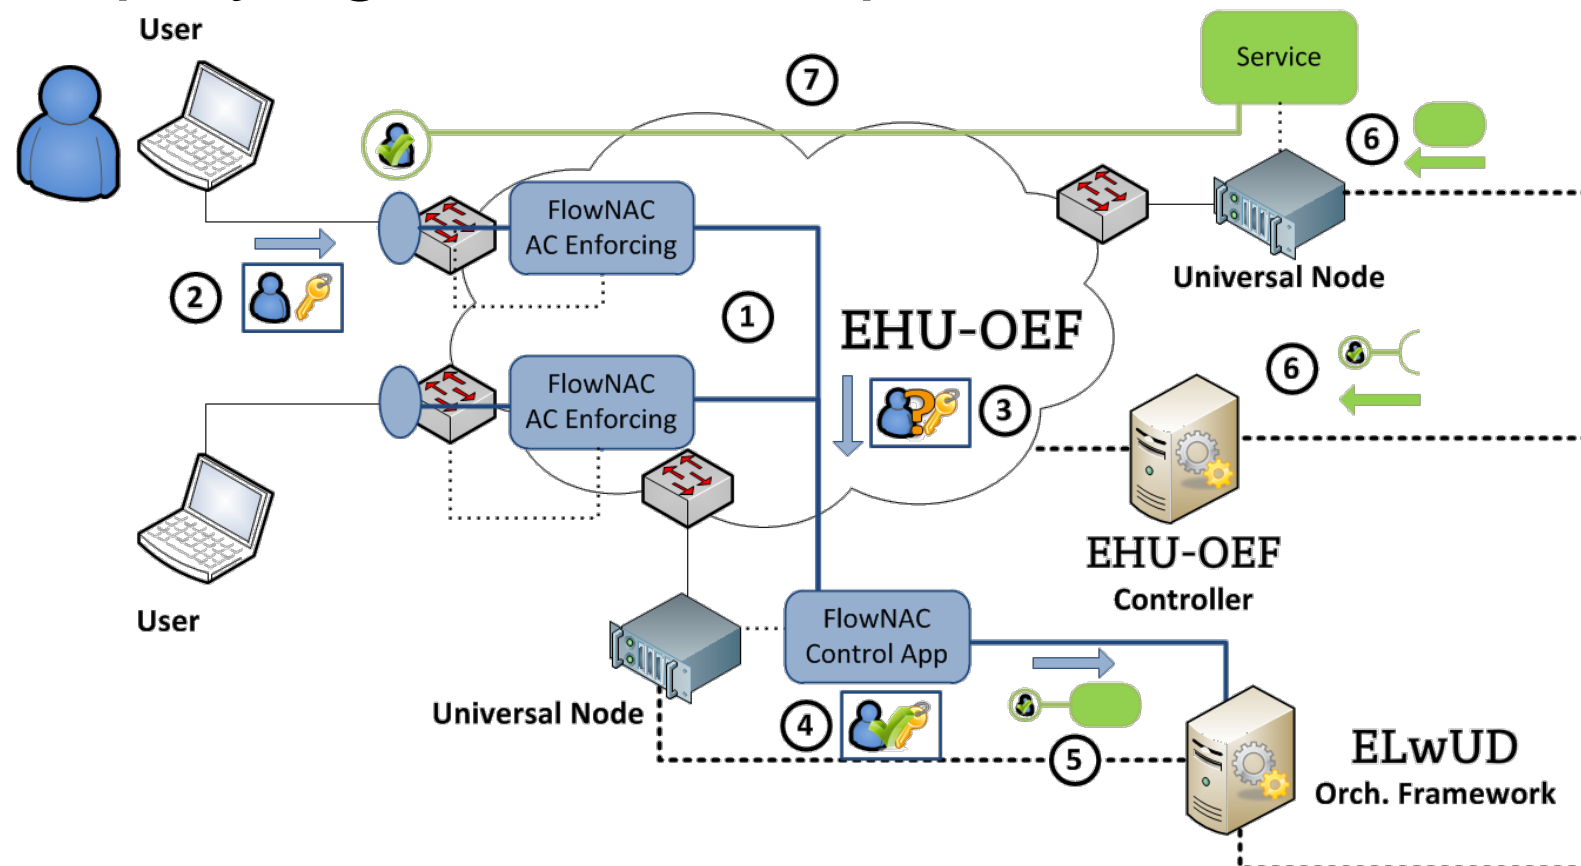
FlowNAC - Demo I

- Access Control VNF use case



FlowNAC - Demo II

- Self-deploying Service Graphs over ELwUD

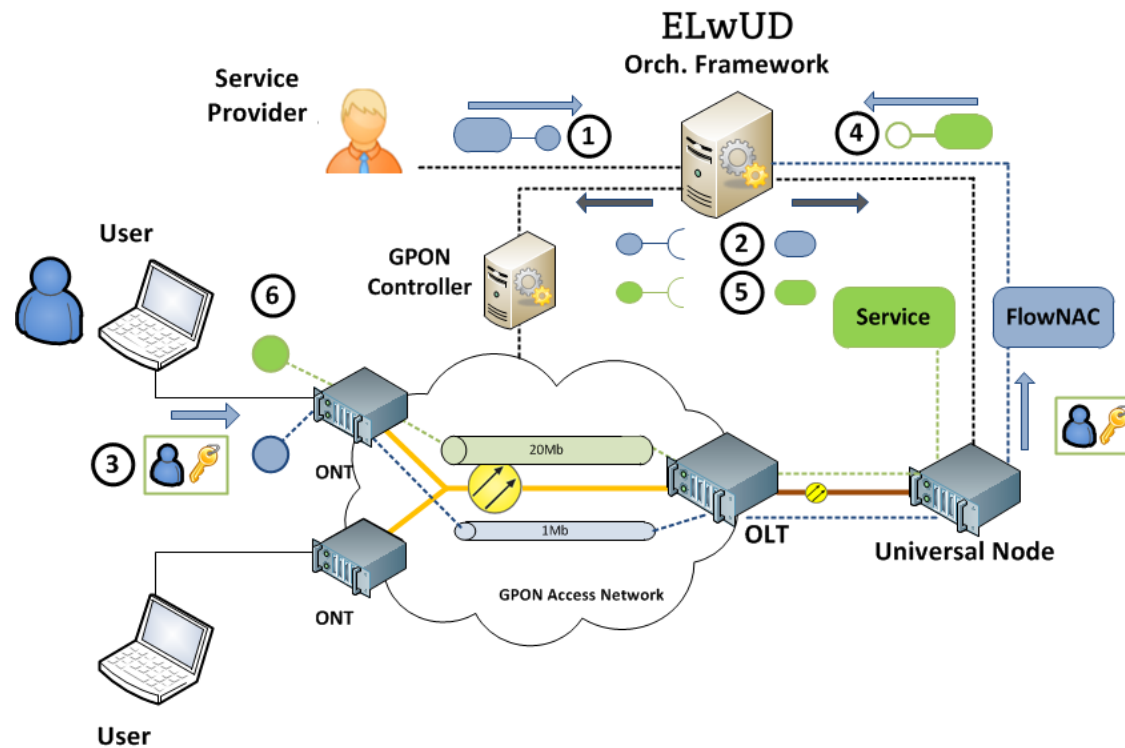


ONS S3 2015
IEEE NetSoft 2015

FlowNAC - Demo III

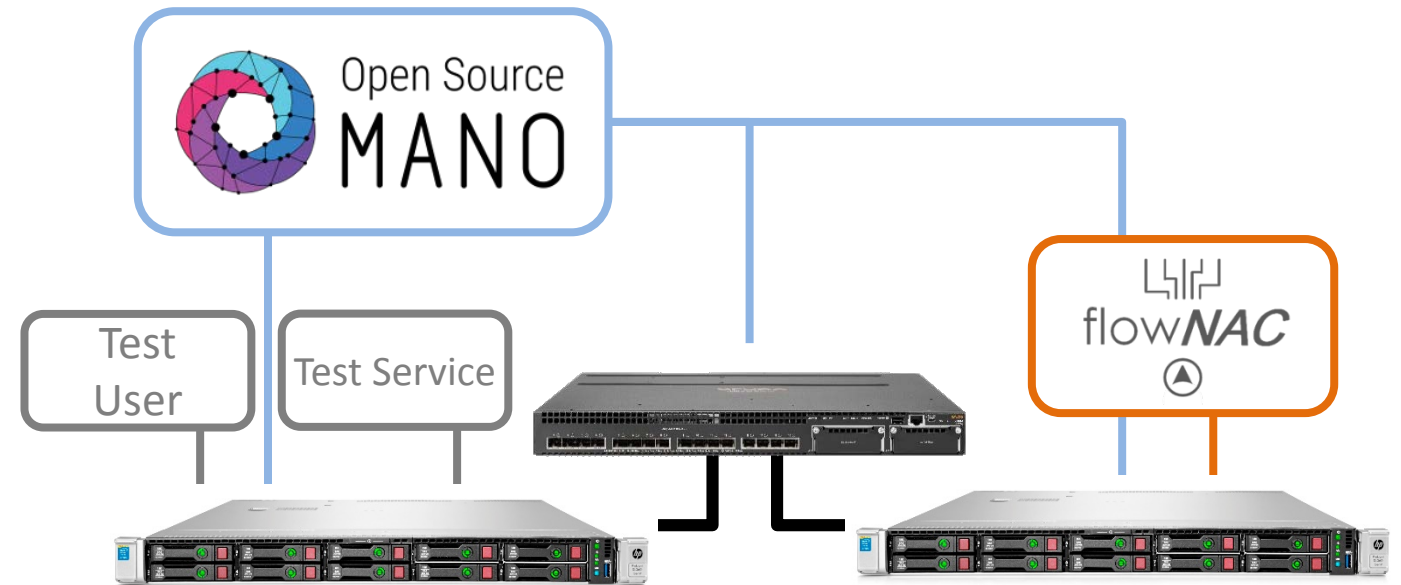
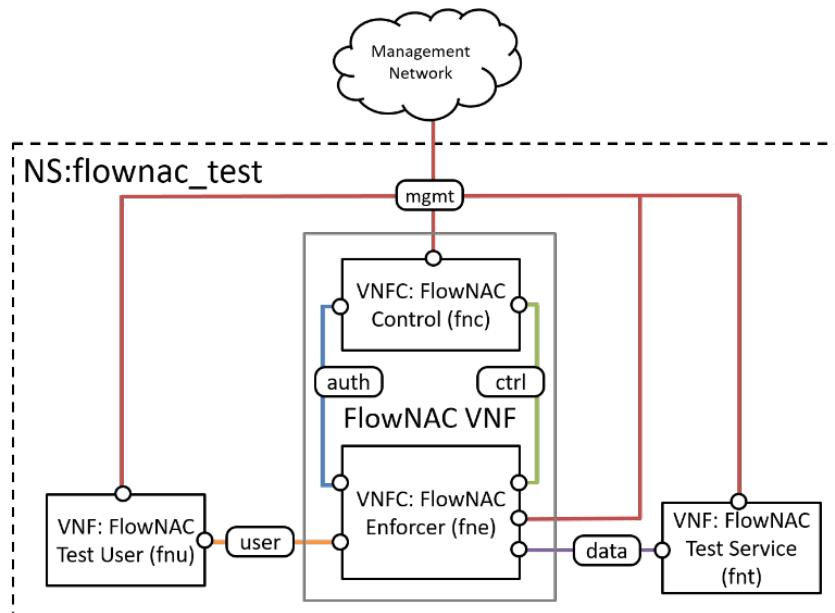
- On demand high speed access over Optical Networks

EWSDN 2015



PoC #43: Scenarios

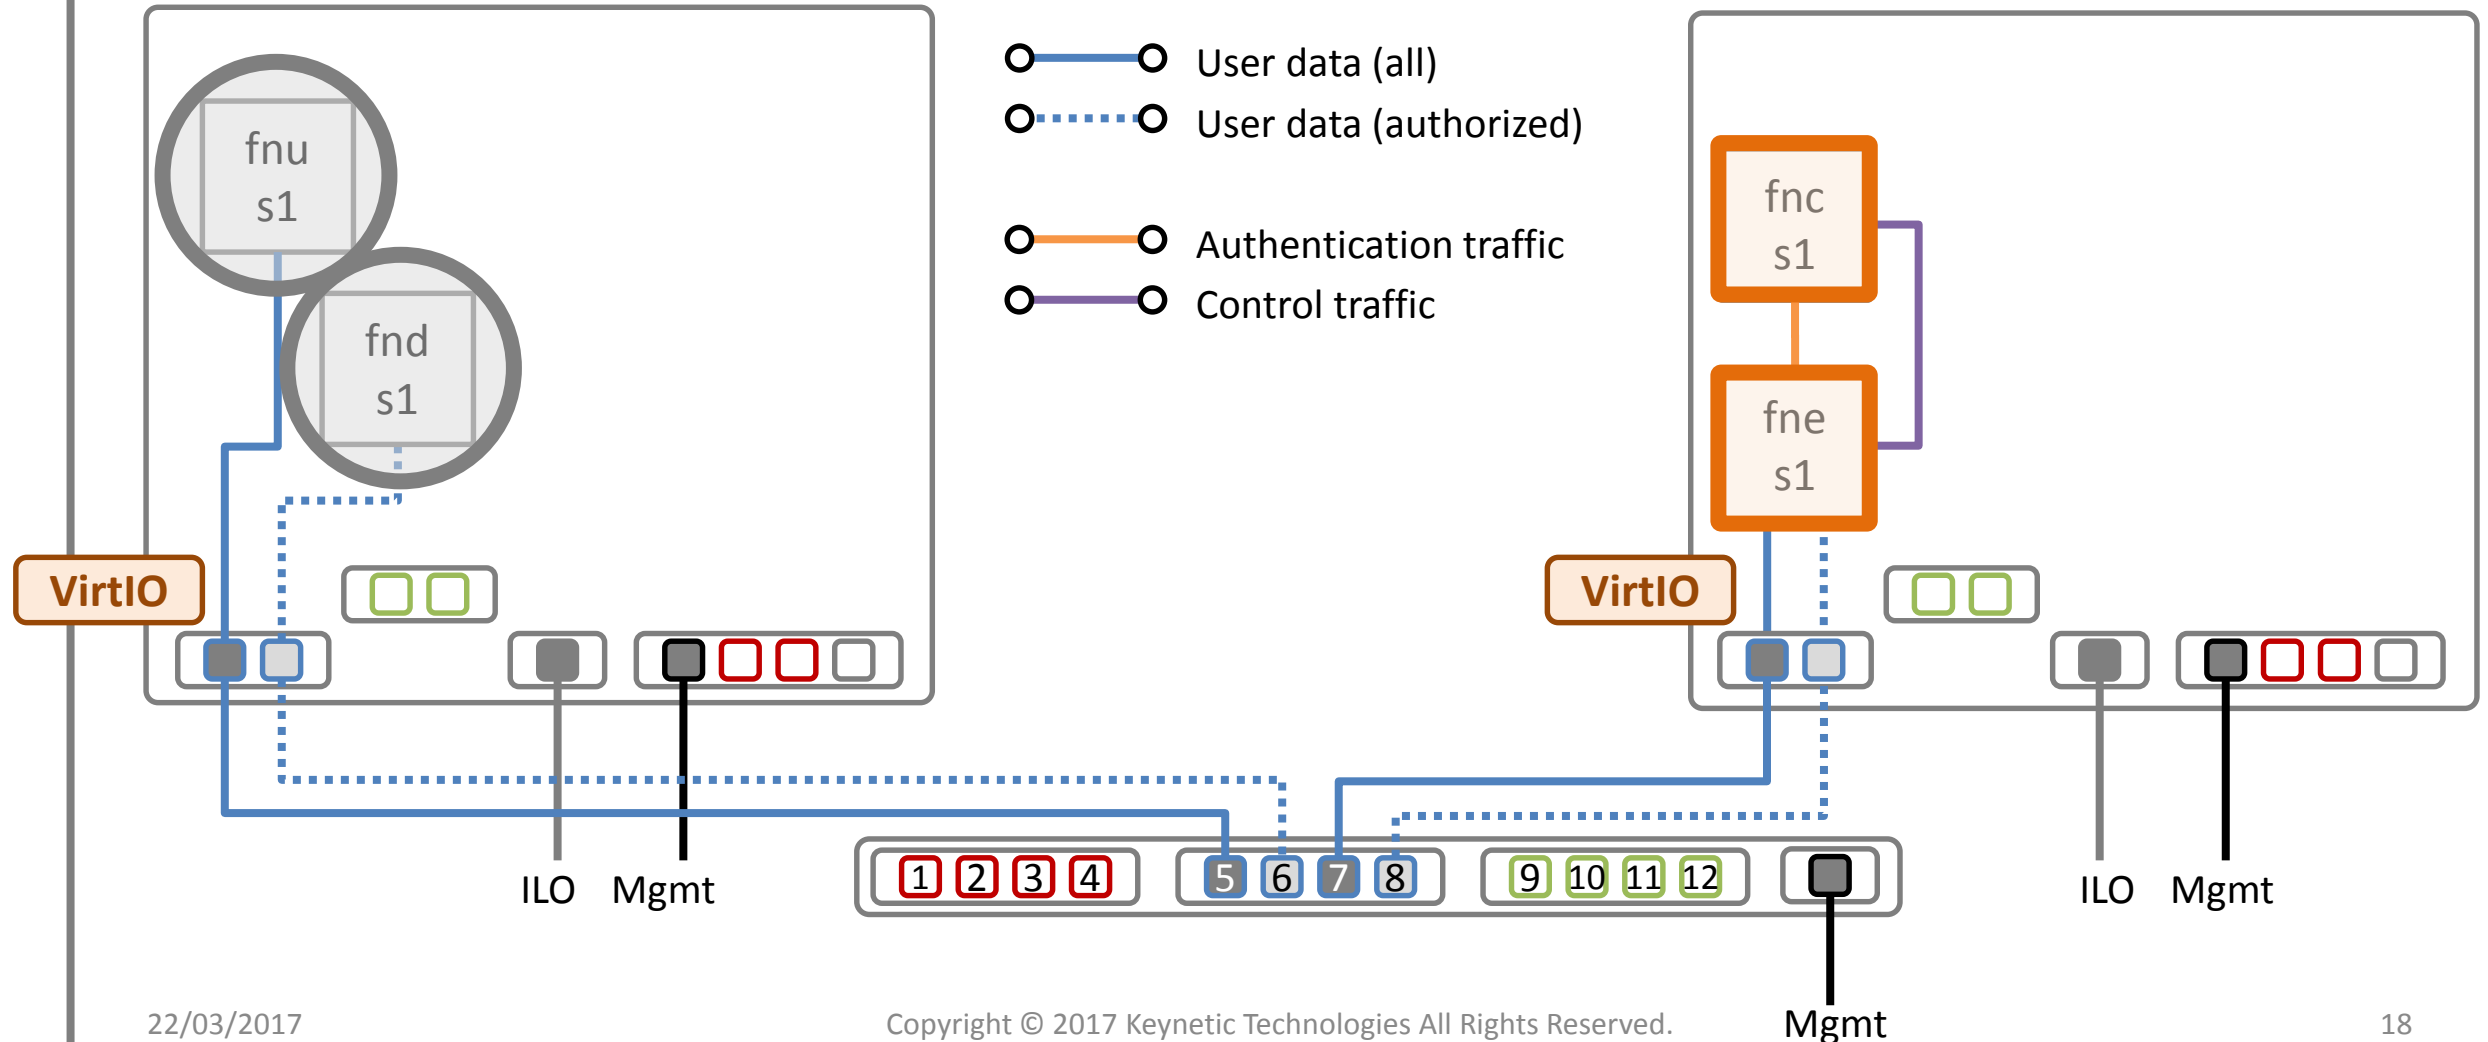
- Scenario 1: Vanilla FNE VNFC (baseline)
- Scenario 2: EPA-aware FNE VNFC
- Scenario 3: SDN-enabled FNE VNFC over SDN fabric (NFVI)



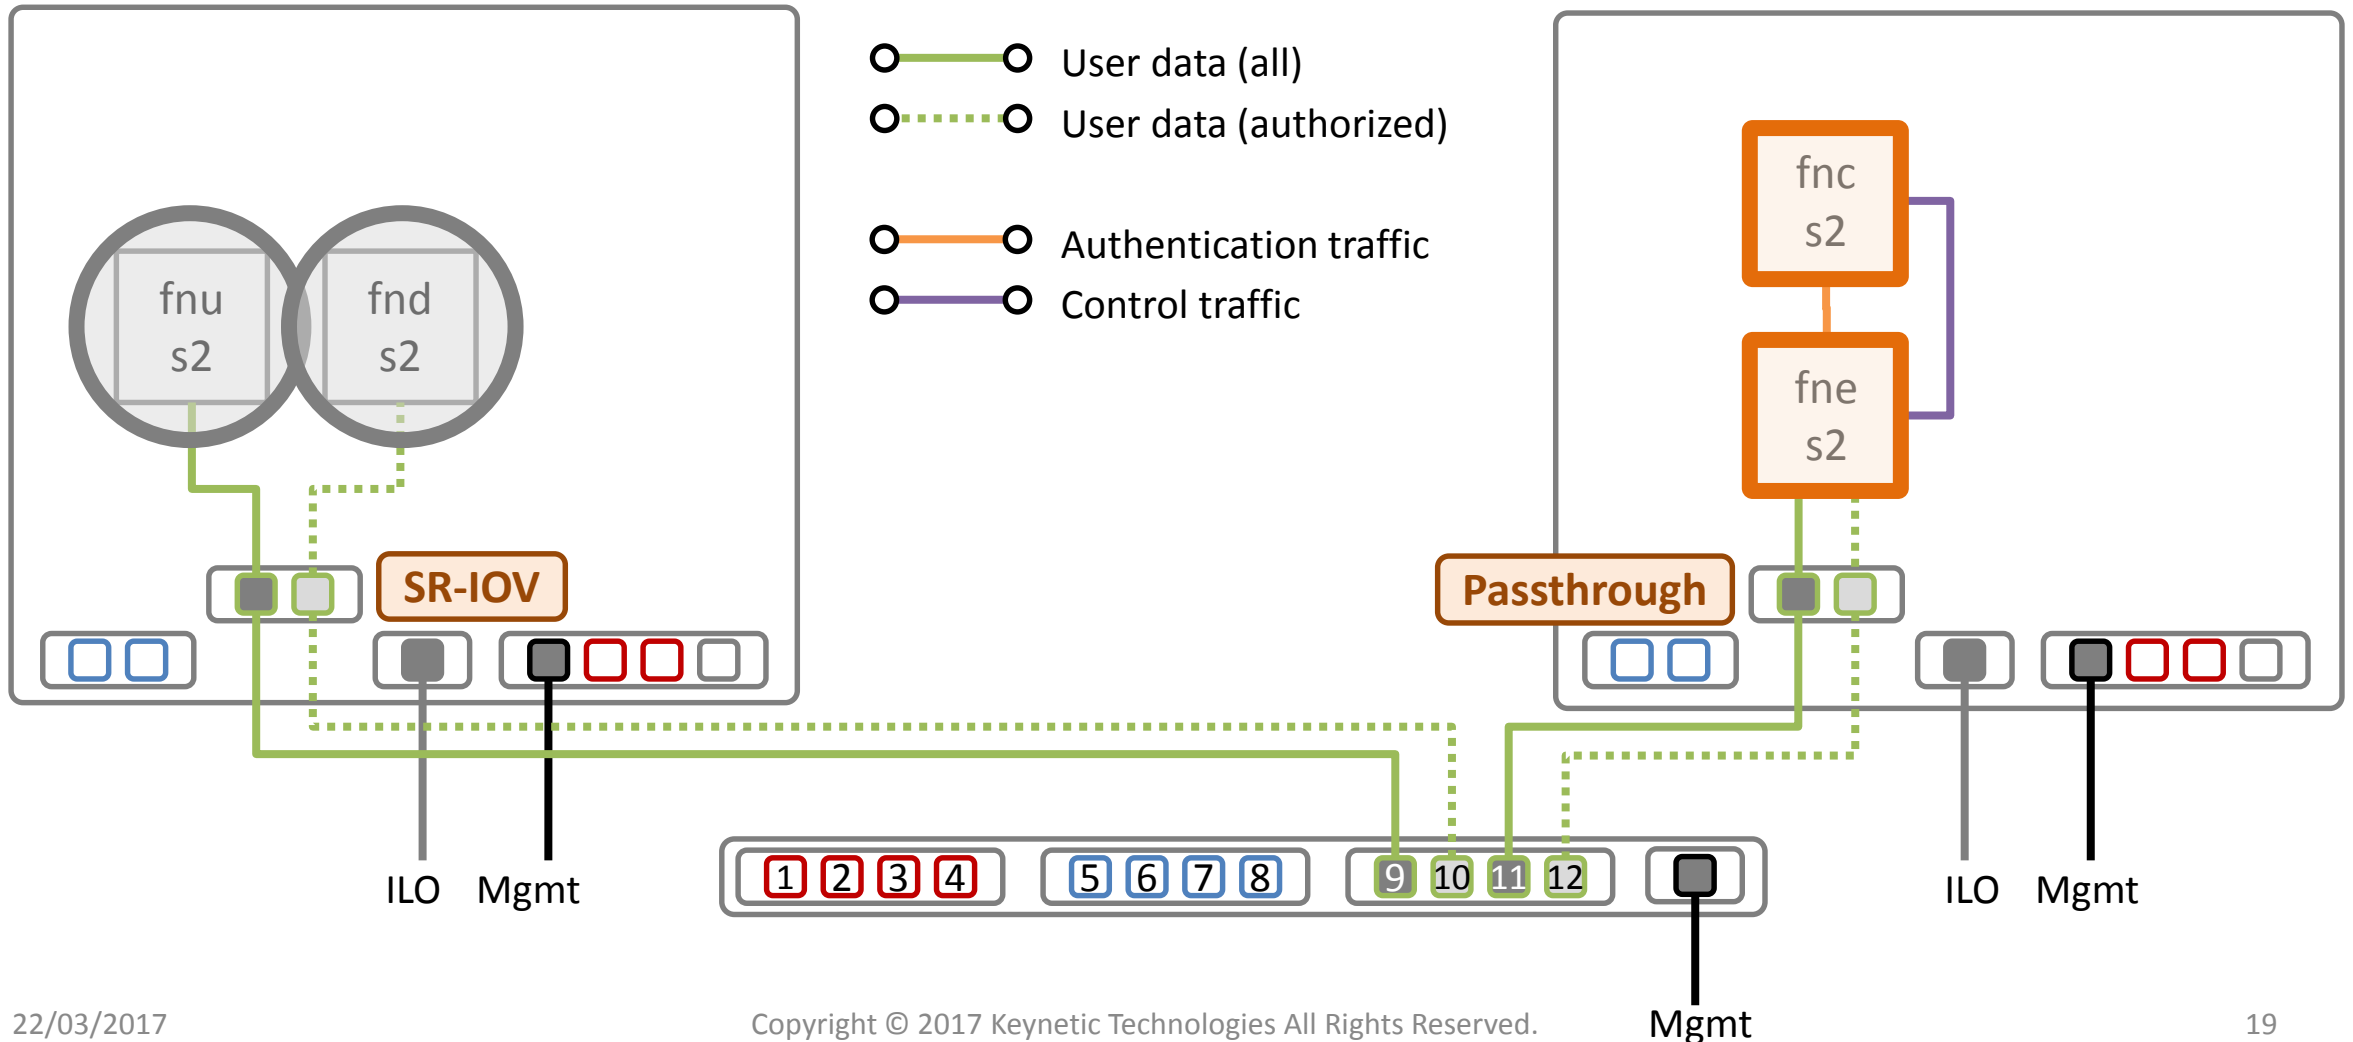
PoC #43: VNF Resources

VNF	Scenario 1 Vanilla	Scenario 2 EPA-aware	Scenario 3 SDN-enabled
FNC	Compute-based, VirtIO networking		
FNE	Compute-based, VirtIO networking	Compute-based, PCI Passthrough networking	SDN-based
FNU FND	Compute-based, VirtIO networking	Compute-based, SR-IOV networking	

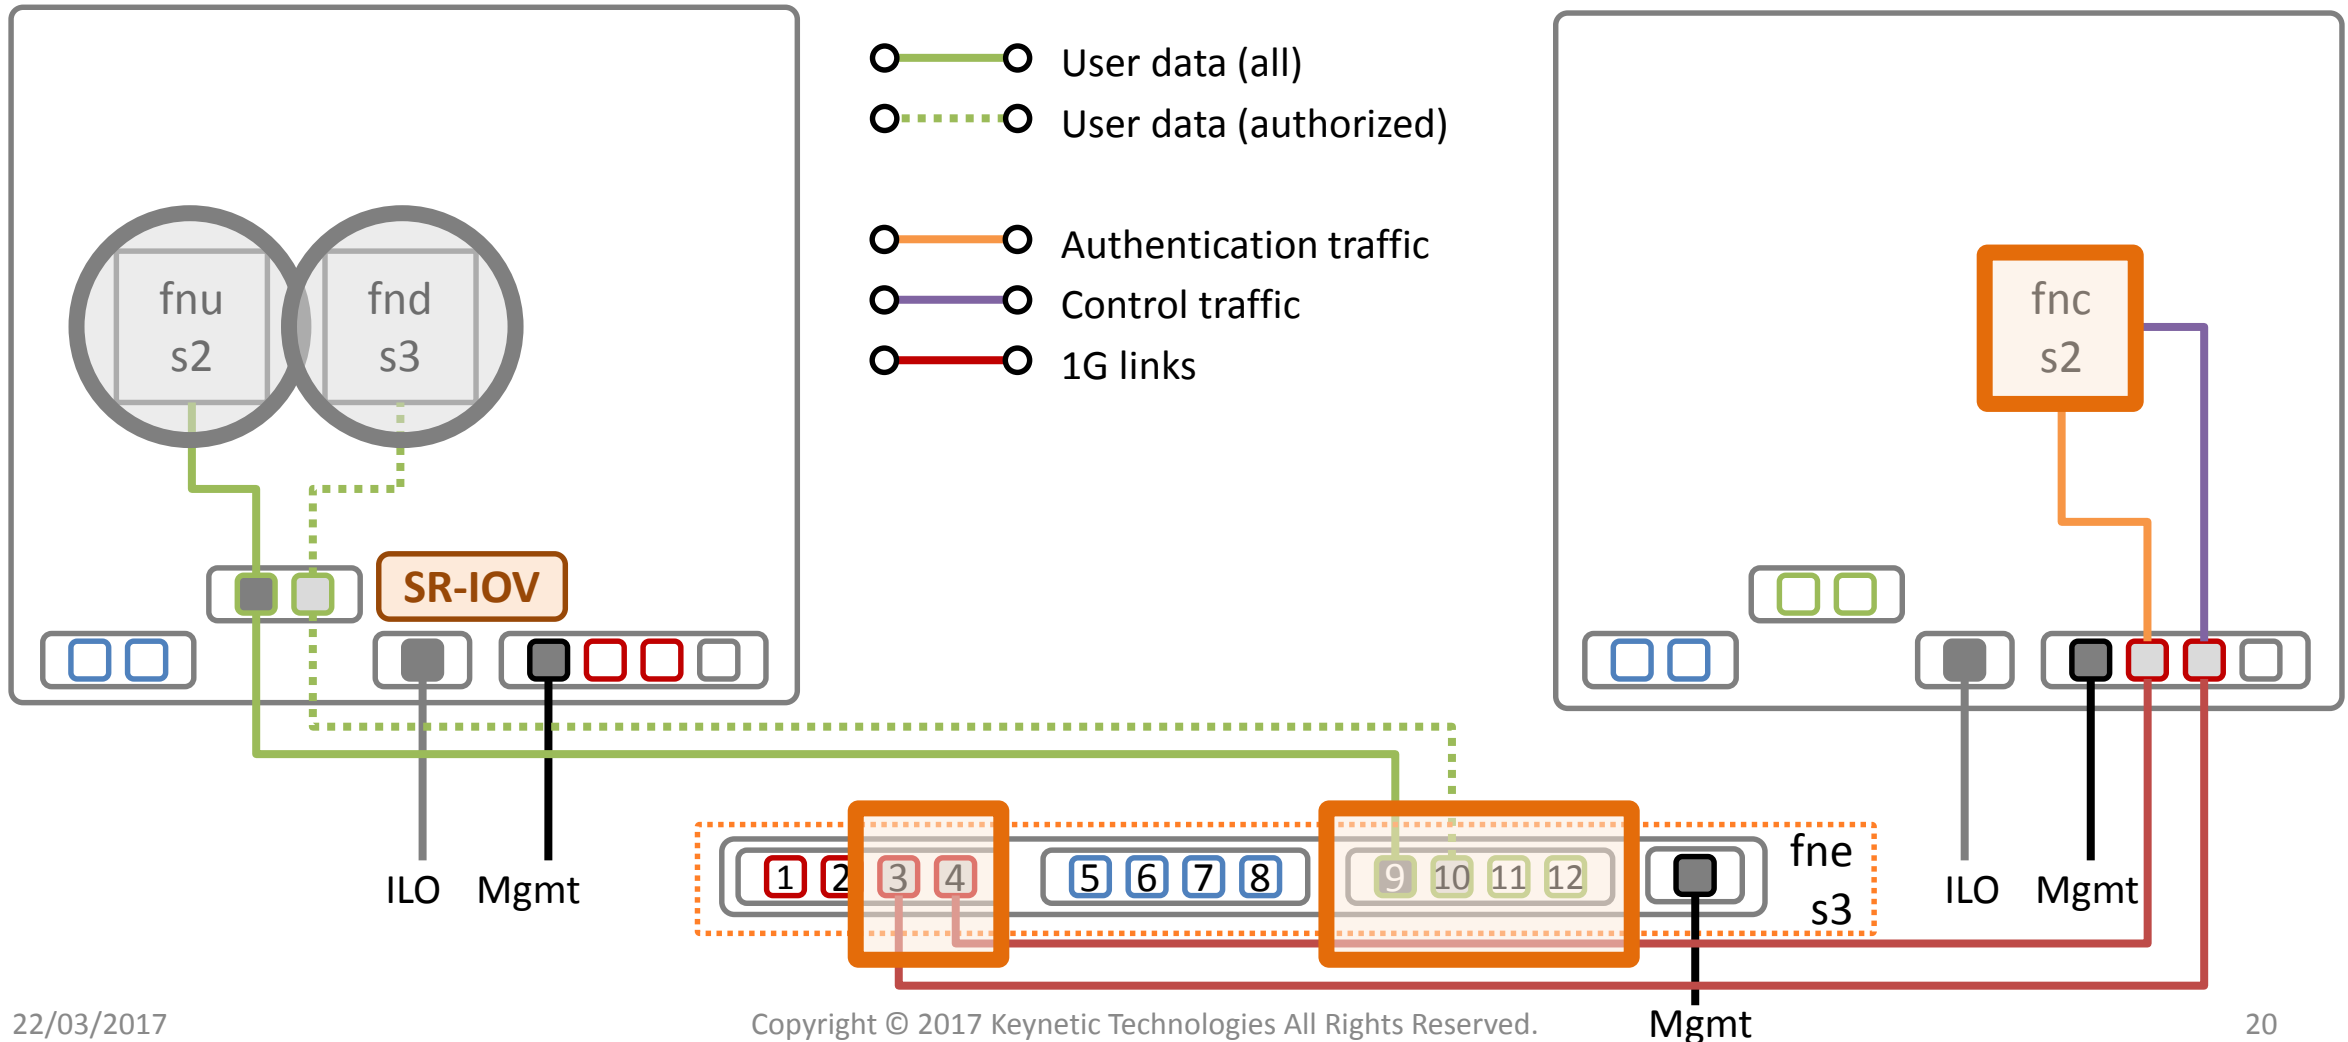
PoC #43: Scenario 1




PoC #43: Scenario 2




PoC #43: Scenario 3




PoC #43: Resources per Scenario



Compute	
Cores	4
Ports	4 (vlan)
Networking	
Ports	4 (vlan)
Traffic	2x 3.5 Gbps *

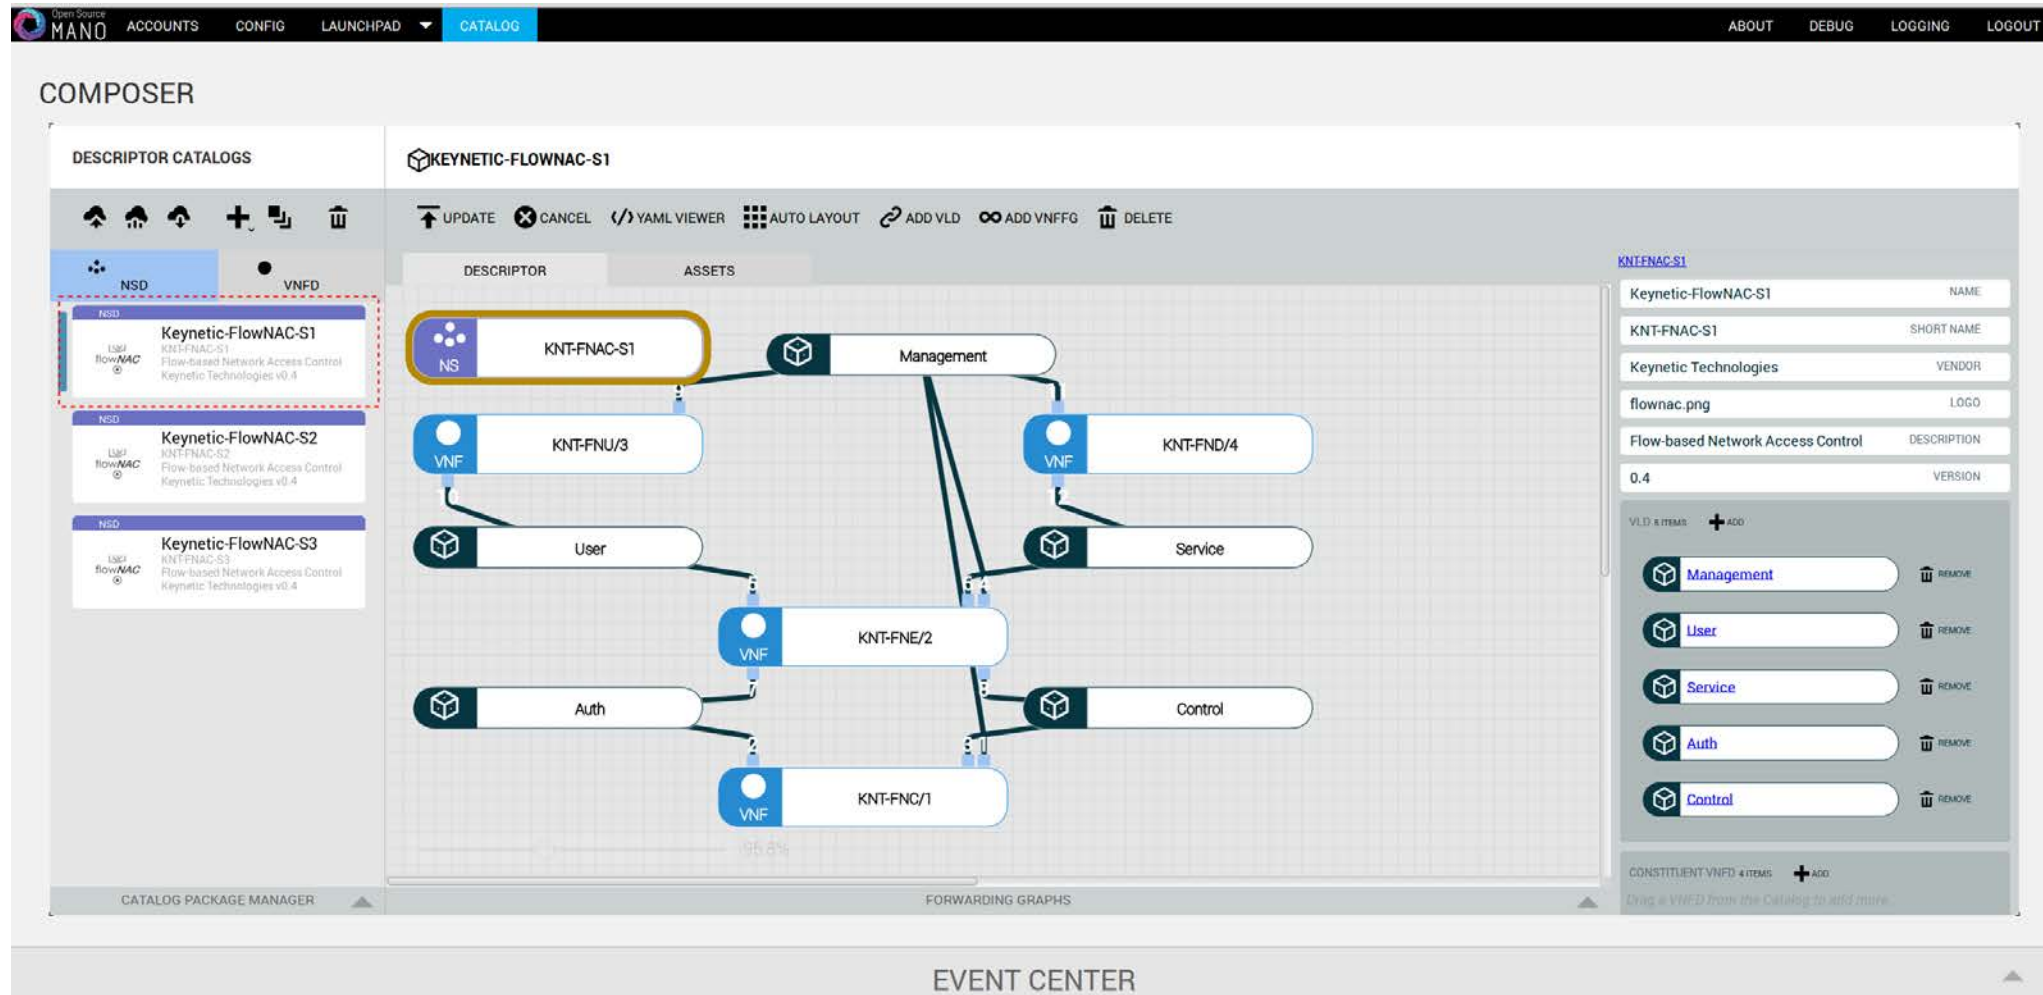


Compute	
Cores	4
Ports	2 (sr-ioiv/vlan) 2 (all)
Networking	
Ports	2 (vlan) + 2 (all)
Traffic	2x 9.8 Gbps *

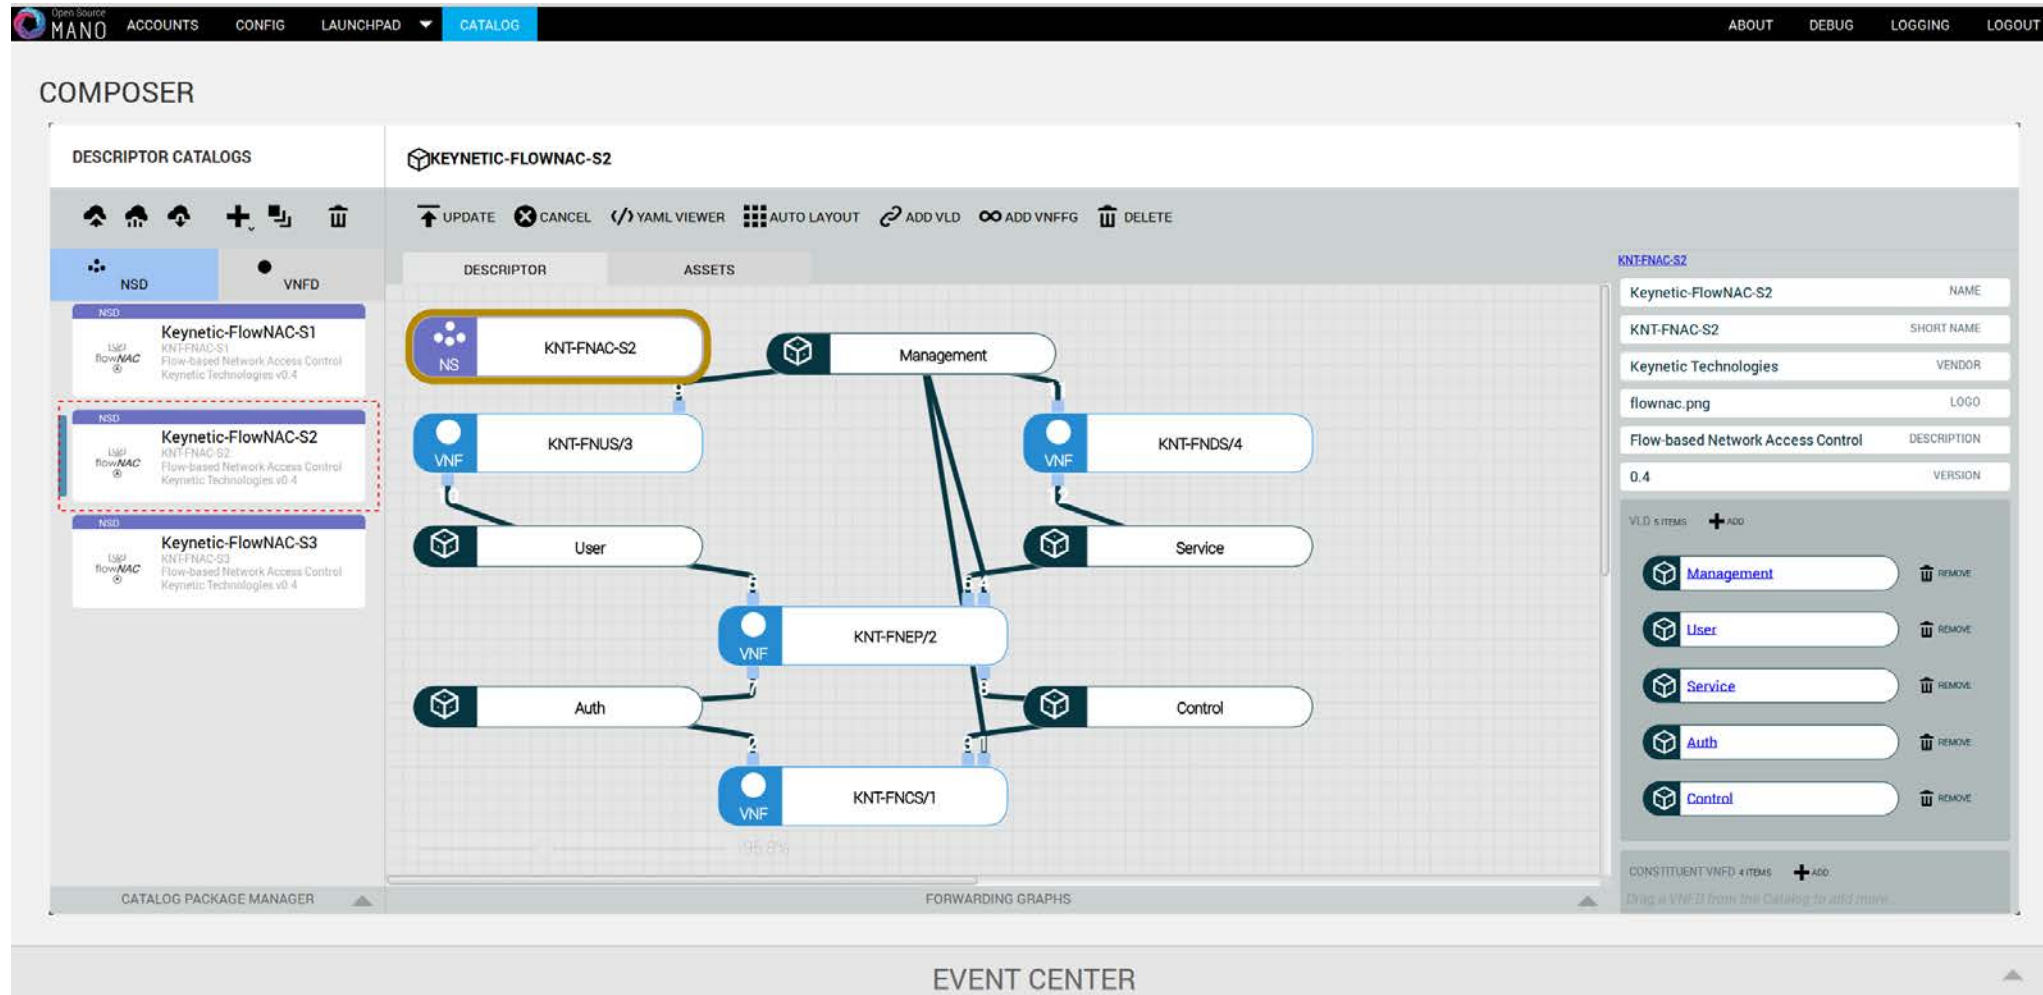


Compute	
Cores	3
Ports	2 (sr-ioiv/vlan) 1 (vlan)
Networking	
Ports	2 (vlan) + 1 (vlan)
Traffic	1x 9.8 Gbps *

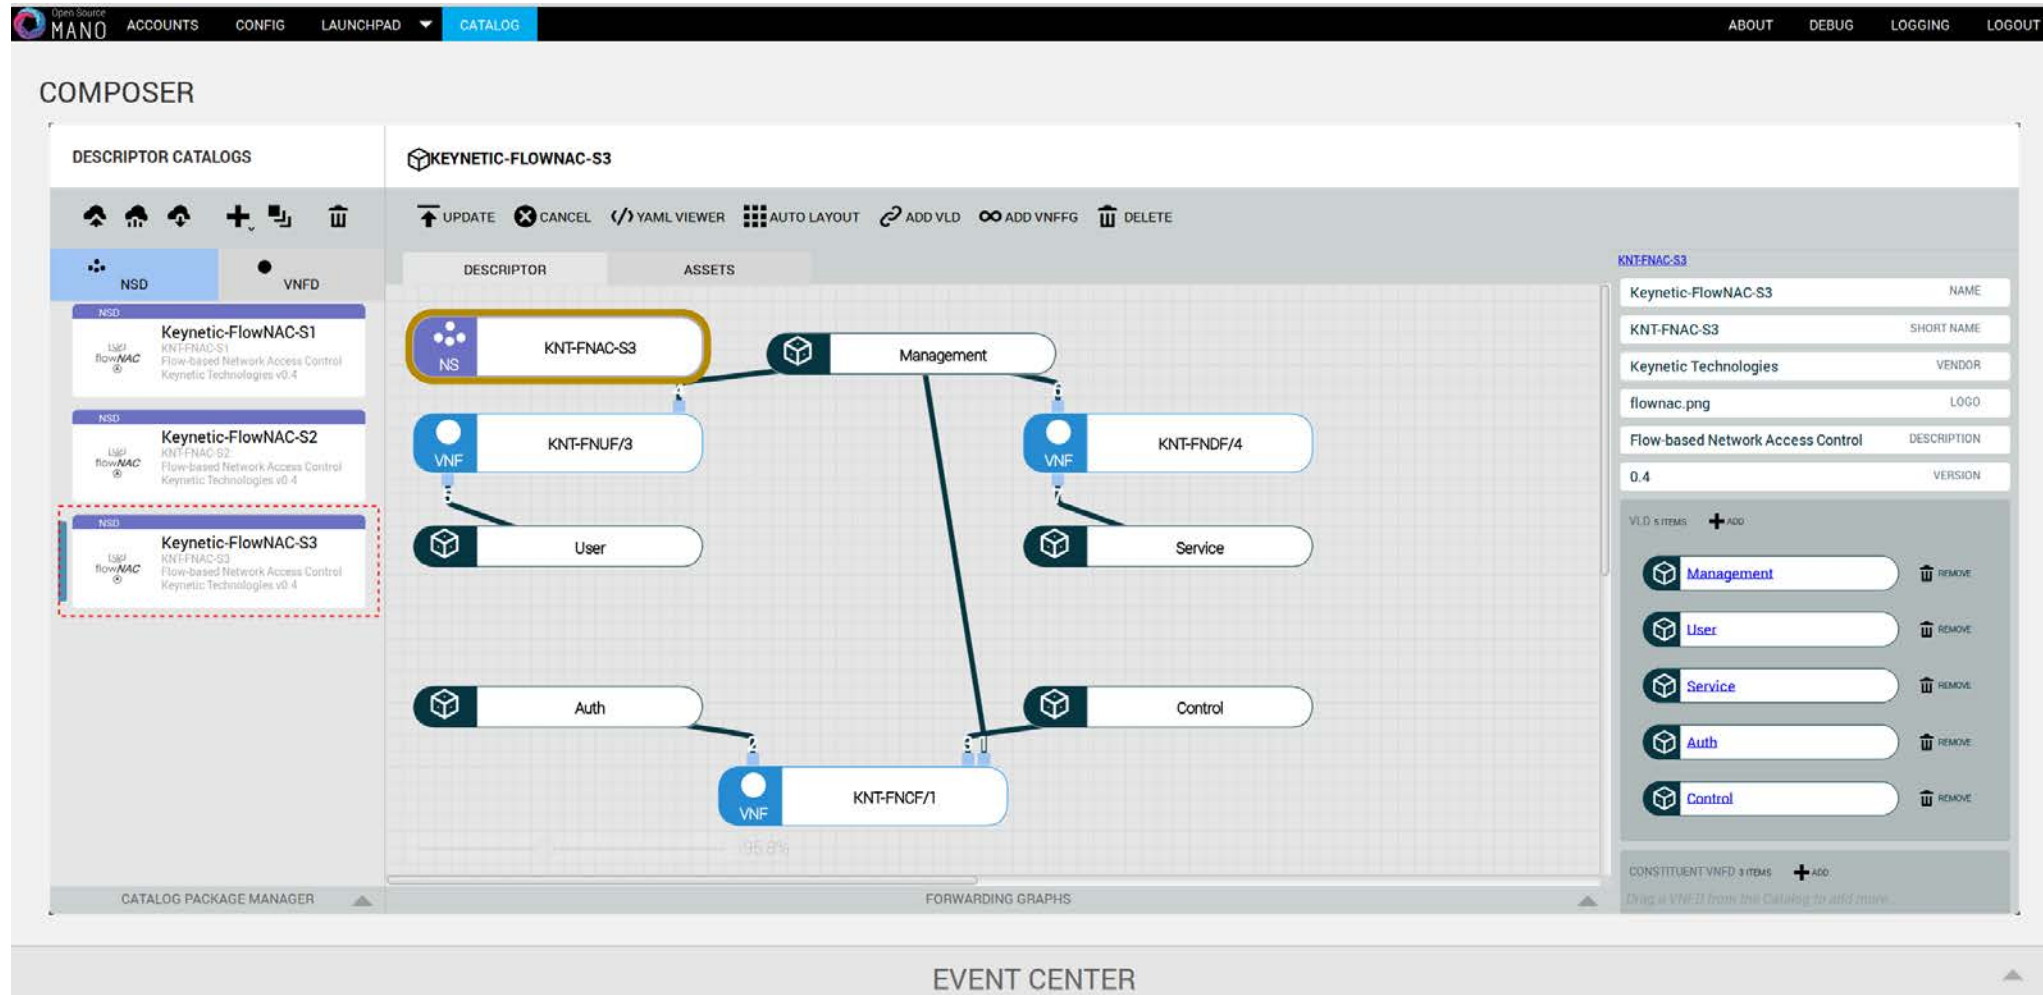
PoC #43: Scenario 1 - NS



PoC #43: Scenario 2 - NS



PoC #43: Scenario 3 - NS



PoC #43: NS Descriptor - Scenario 1

- Mostly a VNF / VL enumeration

Descriptors in OSM repo
under Apache 2 license

```
1 |
2 | # Copyright 2016 Keynetic Technologies
3 | #
4 | # Licensed under the Apache License, Version 2.0 (the "License");
5 | # you may not use this file except in compliance with the License.
6 | # You may obtain a copy of the License at
7 | #
8 | # http://www.apache.org/licenses/LICENSE-2.0
9 | #
10 | # Unless required by applicable law or agreed to in writing, software
11 | # distributed under the License is distributed on an "AS IS" BASIS,
12 | # WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
13 | # See the License for the specific language governing permissions and
14 | # limitations under the License.
15 | #
16 | #
17 | nsd:nsd-catalog:
18 |   nsd:nsd:
19 |     - nsd:vendor: Keynetic Technologies
20 |       nsd:version: '0.4'
21 |       nsd:description: Flow-based Network Access Control
22 |       nsd:id: flownac-s1
23 |       nsd:name: Keynetic-FlowNAC-S1
24 |       nsd:short-name: KNT-FNAC-S1
25 |       nsd:logo: flownac.png
26 |       nsd:constituent-vnfd:
27 |         - nsd:member-vnf-index: '1'
28 |           nsd:start-by-default: 'true'
29 |           nsd:vnfd-id-ref: fnc
30 |         - nsd:member-vnf-index: '2'
31 |           nsd:start-by-default: 'true'
32 |           nsd:vnfd-id-ref: fne
33 |         - nsd:member-vnf-index: '3'
34 |           nsd:start-by-default: 'true'
35 |           nsd:vnfd-id-ref: fnu
36 |         - nsd:member-vnf-index: '4'
37 |           nsd:start-by-default: 'true'
38 |           nsd:vnfd-id-ref: fnd
39 |       nsd:vnfd:
40 |         - nsd:description: Management Link
41 |           nsd:id: 0mgmt
42 |           nsd:mgmt-network: 'true'
43 |           nsd:name: 0mgmt
```

Constituent
VNFs

```
39 | nsd:vnfd:
40 |   nsd:description: Management Link
41 |   nsd:id: 0mgmt
42 |   nsd:mgmt-network: 'true'
43 |   nsd:name: 0mgmt
44 |   nsd:provider-network:
45 |     nsd:overlay-type: VLAN
46 |   nsd:short-name: Management
47 |   nsd:type: ELAN
48 |   nsd:vim-network-name: 0mgmt-net
49 |   nsd:vnfd-connection-point-ref:
50 |     - nsd:member-vnf-index-ref: '1'
51 |       nsd:vnfd-connection-point-ref: 0mgmt
52 |       nsd:vnfd-id-ref: fnc
53 |     - nsd:member-vnf-index-ref: '2'
54 |       nsd:vnfd-connection-point-ref: 0mgmt
55 |       nsd:vnfd-id-ref: fne
56 |     - nsd:member-vnf-index-ref: '3'
57 |       nsd:vnfd-connection-point-ref: 0mgmt
58 |       nsd:vnfd-id-ref: fnu
59 |     - nsd:member-vnf-index-ref: '4'
60 |       nsd:vnfd-connection-point-ref: 0mgmt
61 |       nsd:vnfd-id-ref: fnd
62 |   nsd:description: User Link
63 |   nsd:id: luser
64 |   nsd:mgmt-network: 'false'
65 |   nsd:name: luser
66 |   nsd:provider-network:
67 |     nsd:overlay-type: VLAN
68 |   nsd:short-name: User
69 |   nsd:type: ELAN
70 |   nsd:vim-network-name: luser-net
71 |   nsd:vnfd-connection-point-ref:
72 |     - nsd:member-vnf-index-ref: '2'
73 |       nsd:vnfd-connection-point-ref: luser
74 |       nsd:vnfd-id-ref: fne
75 |     - nsd:member-vnf-index-ref: '3'
76 |       nsd:vnfd-connection-point-ref: luser
77 |       nsd:vnfd-id-ref: fnu
```

Constituent
VLs

```
78 | nsd:description: Service Link
79 | nsd:id: 2data
80 | nsd:mgmt-network: 'false'
81 | nsd:name: 2service
82 | nsd:provider-network:
83 |   nsd:overlay-type: VLAN
84 |   nsd:short-name: Service
85 |   nsd:type: ELAN
86 |   nsd:vim-network-name: 2data-net
87 |   nsd:vnfd-connection-point-ref:
88 |     - nsd:member-vnf-index-ref: '2'
89 |       nsd:vnfd-connection-point-ref: 2data
90 |       nsd:vnfd-id-ref: fne
91 |     - nsd:member-vnf-index-ref: '4'
92 |       nsd:vnfd-connection-point-ref: 2data
93 |       nsd:vnfd-id-ref: fnd
94 |   nsd:description: Auth Link
95 |   nsd:id: 3auth
96 |   nsd:mgmt-network: 'false'
97 |   nsd:name: 3auth
98 |   nsd:provider-network:
99 |     nsd:overlay-type: VLAN
100 |   nsd:short-name: Auth
101 |   nsd:type: ELAN
102 |   nsd:vim-network-name: 3auth-net
103 |   nsd:vnfd-connection-point-ref:
104 |     - nsd:member-vnf-index-ref: '1'
105 |       nsd:vnfd-connection-point-ref: 3auth
106 |       nsd:vnfd-id-ref: fnc
107 |     - nsd:member-vnf-index-ref: '2'
108 |       nsd:vnfd-connection-point-ref: 3auth
109 |       nsd:vnfd-id-ref: fne
110 |   nsd:description: Control Link
111 |   nsd:id: 4ctrl
112 |   nsd:mgmt-network: 'false'
113 |   nsd:name: 4ctrl
114 |   nsd:provider-network:
115 |     nsd:overlay-type: VLAN
116 |     nsd:short-name: Control
117 |     nsd:type: ELAN
118 |     nsd:vim-network-name: 4ctrl-net
119 |     nsd:vnfd-connection-point-ref:
120 |       - nsd:member-vnf-index-ref: '1'
121 |         nsd:vnfd-connection-point-ref: 4ctrl
122 |         nsd:vnfd-id-ref: fnc
123 |       - nsd:member-vnf-index-ref: '2'
124 |         nsd:vnfd-connection-point-ref: 4ctrl
125 |         nsd:vnfd-id-ref: fne
126 |
```

PoC #43: VNF Descriptors - FNC

- Same descriptor all cases

Header info

```
17 vnfd:vnfd-catalog:
18   vnfd:
19     vendor: Keynetic Technologies
20     version: '0.4'
21     description: FlowNAC Control
22     id: fnc
23     name: Keynetic-FNC
24     short-name: KNT-FNC
25     logo: flownac.png
26     service-function-chain: UNMANAGED
27     connection-point:
28       - id: 0mgmt
29         name: 0mgmt
30         type: VPORT
31       - id: 3auth
32         name: 3auth
33         type: VPORT
34       - id: 4ctrl
35         name: 4ctrl
36         type: VPORT
```

Connection
points

Interfaces

```
37   vdu:
38     - description: FNC VDU
39       external-interface:
40         - name: 0eth0
41           virtual-interface:
42             type: OM-MGMT
43             vnfd-connection-point-ref: 0mgmt
44         - name: 3auth
45           virtual-interface:
46             type: VIRTIO
47             vnfd-connection-point-ref: 3auth
48         - name: 4ctrl
49           virtual-interface:
50             type: VIRTIO
51             vnfd-connection-point-ref: 4ctrl
```

EPA
requirements

```
52   guest-epa:
53     cpu-pinning-policy: SHARED
54     cpu-thread-pinning-policy: ISOLATE
55     mempage-size: LARGE
56     numa-node-policy:
57       mem-policy: STRICT
```

Image
requirements

```
58   id: fnc-vdu
59   image: fnc
60   name: fnc-vdu
61   vm-flavor:
62     memory-mb: '2048'
63     storage-gb: '4'
64     vcpu-count: '2'
```

Configuration
primitives

```
65   vnf-configuration:
66     config-attributes:
67       config-delay: 10
68     service-primitive:
69       - name: config
70         parameter:
71           - name: ssh-hostname
72             data-type: STRING
73             mandatory: 'true'
74           - name: ssh-username
75             data-type: STRING
76             mandatory: 'true'
77           - name: ssh-password
78             data-type: STRING
79             mandatory: 'true'
80           - name: mode
81             data-type: STRING
82             read-only: 'true'
83             default-value: fnc
84       - name: start
85       - name: stop
86       - name: restart
```

Initial
configuration

```
87   initial-config-primitive:
88     - name: config
89       parameter:
90         - name: ssh-hostname
91           value: <rw_mgmt_ip>
92         - name: ssh-username
93           value: flownac
94         - name: ssh-password
95           value: fnac@osm
96         - name: mode
97           value: fnc
98         seq: '1'
99       - name: start
100         seq: '2'
101   juju:
102     charm: flownac
```


PoC #43: VNF Descriptors - FNE

- Different descriptors in Scenario 1 and Scenario 2
 - Change virtual interface type: VirtIO → Passthrough

VirtIO
interfaces

```
43 vdu:
  type: FNE VDU
  interface:
    name: 0eth0
    type: OM-MGMT
  vnfd-connection-point-ref: 0mgmt
49 - name: 1user
50   virtual-interface:
51     type: VIRTIO
52     bandwidth: '10000000000'
53   vnfd-connection-point-ref: 1user
54 - name: 2data
55   virtual-interface:
56     type: VIRTIO
57     bandwidth: '10000000000'
58   vnfd-connection-point-ref: 2data
59 - name: 3auth
60   virtual-interface:
61     type: VIRTIO
62   vnfd-connection-point-ref: 3auth
63 - name: 4ctrl
64   virtual-interface:
65     type: VIRTIO
66   vnfd-connection-point-ref: 4ctrl
67
```

Passthrough
interfaces

```
43 vdu:
  type: FNEP VDU
  interface:
    name: 0eth0
    type: OM-MGMT
  vnfd-connection-point-ref: 0mgmt
49 - name: 1user
50   virtual-interface:
51     type: PCI-PASSTHROUGH
52     vpci: '0000:00:10.0'
53     bandwidth: '10000000000'
54   vnfd-connection-point-ref: 1user
55 - name: 2data
56   virtual-interface:
57     type: PCI-PASSTHROUGH
58     vpci: '0000:00:11.0'
59     bandwidth: '10000000000'
60   vnfd-connection-point-ref: 2data
61 - name: 3auth
62   virtual-interface:
63     type: VIRTIO
64   vnfd-connection-point-ref: 3auth
65 - name: 4ctrl
66   virtual-interface:
67     type: VIRTIO
68   vnfd-connection-point-ref: 4ctrl
69
```

Scenario 3 not supported in
OSM Release 1

```
17 vnfd:vnfd-catalog:
18   vnfd:
19     - vendor: Keynetic Technologies
20       version: '0.4'
21       description: FlowNAC Enforcement
22       id: fne
23       name: Keynetic-FNE
24       short-name: KNT-FNE
25       logo: flownac.png
26       service-function:
27         connection-point-ref: 1user
28         - id: 1user
29           name: 1user
30           type: VPORT
31         - id: 2data
32           name: 2data
33           type: VPORT
34         - id: 3auth
35           name: 3auth
36           type: VPORT
37         - id: 4ctrl
38           name: 4ctrl
39           type: VPORT
40         - id: 4ctrl
41           name: 4ctrl
42           type: VPORT
43   vdu:
44     - name: 1user
45       type: VIRTIO
46       bandwidth: '10000000000'
47       vnfd-connection-point-ref: 1user
48     - name: 2data
49       type: VIRTIO
50       bandwidth: '10000000000'
51       vnfd-connection-point-ref: 2data
52     - name: 3auth
53       type: VIRTIO
54       bandwidth: '10000000000'
55       vnfd-connection-point-ref: 3auth
56     - name: 4ctrl
57       type: VIRTIO
58       bandwidth: '10000000000'
59       vnfd-connection-point-ref: 4ctrl
60
```

PoC #43: VNF Descriptors - FN{U|D}

- Descriptors in Scenario 1 and Scenarios 2/3
 - Change virtual interface type: VirtIO → SR-IOV

```
34 vdu:
35   - description: FNU VDU
36     external-interface:
37       - name: 0eth0
38         virtual-interface:
39           type: OM-MGMT
40         vnfd-connection-point-ref: 0mgmt
41       - name: luser
42         virtual-interface:
43           type: VIRTIO
44           bandwidth: '10000000000'
45           vnfd-connection-point-ref: luser
46     guest-epa:
47       cpu-pinning-policy: SHARED
48       cpu-thread-pinning-policy: ISOLATE
49       mempage-size: LARGE
50       numa-node-policy:
51         mem-policy: STRICT
52     id: fnu-vdu
53     image: fnu
54     name: fnu-vdu
55     vm-flavor:
56       memory-mb: '2048'
57       storage-gb: '4'
58       vcpu-count: '2'
```

VirtIO interface

```
34 vdu:
35   - description: FNUS VDU
36     external-interface:
37       - name: eth0
38         virtual-interface:
39           type: OM-MGMT
40         vnfd-connection-point-ref: 0mgmt
41       - name: user
42         virtual-interface:
43           type: SR-IOV
44           vpci: '0000:05:10.0'
45           bandwidth: '10000000000'
46           vnfd-connection-point-ref: luser
47     guest-epa:
48       cpu-pinning-policy: SHARED
49       cpu-thread-pinning-policy: ISOLATE
50       mempage-size: LARGE
51       numa-node-policy:
52         mem-policy: STRICT
53     id: fnus-vdu
54     image: fnus
55     name: fnus-vdu
56     supplemental-boot-data:
57       boot-data-drive: 'false'
58     vm-flavor:
59       memory-mb: '2048'
60       storage-gb: '4'
61       vcpu-count: '2'
```

SR-IOV interface

PoC #43: Process

Step	Scenario 1 Vanilla	Scenario 2 EPA-aware	Scenario 3 SDN-enabled
Instantiate	<p>Instantiate NS from OSM</p> <ul style="list-style-type: none">• Provision VNFs• Provision paths (VLANs)• Initial VNF configuration		<p>Very similar to what's done for the compute-based VNFs</p>
Operate	<p>Launch actions through configuration channel</p> <ul style="list-style-type: none">• Check access before authentication• Authenticate• Measure resource usage and performance		

PoC #43: Initial learnings & Next steps



- Resources efficiency
 - Less cores, less backplane, less (dedicated) ports
- Bump-in-the-wire requirements
 - Careful with VIM networking ‘interference’
 - Passthrough (dedicated ports)
- Interfaces not ready for networking resources
 - (Currently, just) vlan pipes
- From an abstract perspective not much difference between provisioning compute and networking resources
 - Still not proper slicing/isolation (both data and control planes)
 - Still not proper abstractions/APIs
- Measure with multiple packet sizes (10 Gbps with small packets)
- Scenario 2 with DPDK
- Measure latency



Keynetic Technologies

<info@keynetic.es>