
A.1 NFV ISG PoC Proposal

A.1.1 PoC Team Members

- PoC Project Name: Towards an efficient Data Plane processing.
- Network Operators / Service Providers:
 - Telefónica (Contact: Francisco Javier Ramón Salguero javier.ramon@telefonica.com)
- Manufacturer A:
 - Keynetic Technologies (Contact: Jokin Garay jokin.garay@keynetic.es)
- Manufacturer B:
 - Hewlett Packard Enterprise (Contact: Enrique Matorras enrique.matorras@hpe.com)
- Additional members (research):
 - University of the Basque Country UPV/EHU (Contact: Eduardo Jacob eduardo.jacob@ehu.eus)

A.1.2 PoC Project Goals

The PoC will explore one fundamental criteria to consider at the VNF design phase to decompose the VNF into separate sub-components: stateful network function components (compute) and stateless data path processing components (networking).

This separation of VNF into sub-components (related to the micro-service approach) will allow each sub-component to independently scale depending on the evolution of each deployment.

The PoC will also explore the improvement of efficiency in VNF data processing while minimizing the overall resources used (resource occupation) from the NFVI, mainly when considering stateless data path processing components.

The PoC will evaluate different alternatives to deploy the stateless data path processing component using both compute and networking resources to implement the functionality. In the latter case, the VNF will rely on SDN technology to implement the functionality. It is worth noting that the PoC is not committed to apply any general mechanism for SDN/NFV integration (not available to the best of our knowledge), nor to explore those general mechanisms, though we expect to derive some useful insights in this respect

The PoC will furthermore propose pointers to ease the integration of stateful data path processing into the ETSI NFV architectural framework.

A.1.3 PoC Demonstration

- Venue for the demonstration of the PoC: NFV#17 meeting, 21-24 February 2017, held at the University of the Basque Country UPV/EHU, Bilbao, Spain.

A.1.4 Publication

- What would be the publication channel(s) for the PoC. Keynetic Technologies' website for the FlowNAC VNF, I2T Research Group website.
- What would be the planned publication date(s)? May 2017

- URLs where applicable: <http://www.flownac.net> (offline for the time being) , <http://i2t.ehu.eus>, <http://www.ehu.eus>

We also intend to use the PoC for public demos at trade-shows and appropriate industry events. A final PoC report will be submitted to ETSI NFV ISG by 31st May, 2017.

A.1.5 PoC Project Timeline

- | | |
|---|--|
| • What is the PoC start date? | Project already started |
| • (First) Demonstration target date | 21-24/02/2017 |
| • PoC stages target dates(optional) | No stages currently planned |
| • PoC Report target date | 31/05/2017 |
| • When is the PoC considered completed? | Once the final PoC report has been published to the ETSI NFV ISG |

A.2 NFV PoC Technical Details

A.2.1 PoC Overview

The PoC provides a use case with a VNF that demonstrates the applicability of the ideas described in the PoC Project Goals section.

In particular, the VNF has already participated in the 1st NFV Plugtest, following the methodology described in NFV-TST002.

In order to validate the concept the PoC is built around a VNF that implements FlowNAC, a function that allows fine-grained control of which traffic from the user is granted access to the network. The users are authenticated and authorized on a per-service basis. All the incoming traffic from the user is independently evaluated and categorized as a specific service (or none) at the data plane. Then a basic allow or deny decision is enforced for each frame depending on the associated service being authorized or not.

The FlowNAC VNF is decomposed in two VNFCs (see Figure 1): the FlowNAC Controller (FNC), and the FlowNAC Enforcer (FNE). On the one hand, the authentication and authorization traffic from the end-users is redirected by the FNE to the assigned FNC, which is responsible of keeping the state of the AAA processes. On the other hand, the FNE is responsible of enforcing the policy (allow or deny access) limiting the access only to already authorized services. This separation within FlowNAC distinguishes between the service control plane (related to the FNC), and the service data plane (related to the FNE).

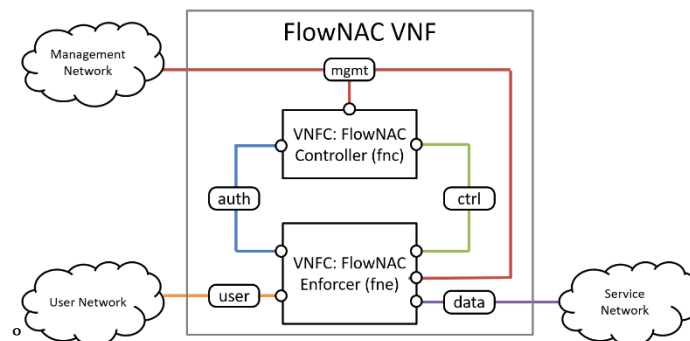


Figure 1 FlowNAC VNF and VNFCs: FlowNAC Controller (FNC) and FlowNAC Enforcer (FNE)

With regard to the implementation, the FNC is implemented with compute resources (e.g. VM), whereas the implementation of the FNE, the focus of this PoC, can use both compute or networking resources.

When the enforcing FNE is implemented as a VM, the traffic coming from the user has to pass through the network infrastructure to reach the assigned compute resource and corresponding VM, where the access control policy will be applied, the authorized traffic will then go down to the network infrastructure to the next VNF in the chain. This means the compute resources are consumed to perform the enforcing functionality. If this FNE functionality is offloaded to the underlying networking resources (that must be anyway traversed), the NFVI saves the compute resources and avoids double processing of authorized traffic by network resources.

The main benefits from the decoupled VNF design are the efficient data processing achieved by optimized hardware, optimization of resources by means of avoiding data traffic going up/down to/from a VM, and independent scalability of each component. This independence is fundamental when the stateless and stateful processing is unbalanced. In this case, the most demanding processing can increase its assigned resources independently and optimize the resource utilization. There are other VNFs that could benefit from this approach, i.e.: load balancer.

A.2.2 PoC Scenarios

The FlowNAC VNF and separation between FNC and FNE will be validated in the following testing setup. The full virtualized setup is shown in Figure 2, in addition to the FlowNAC VNF, two test VNFs are deployed: the FlowNAC Test User (FNU) and the FlowNAC Test Service (FNT).

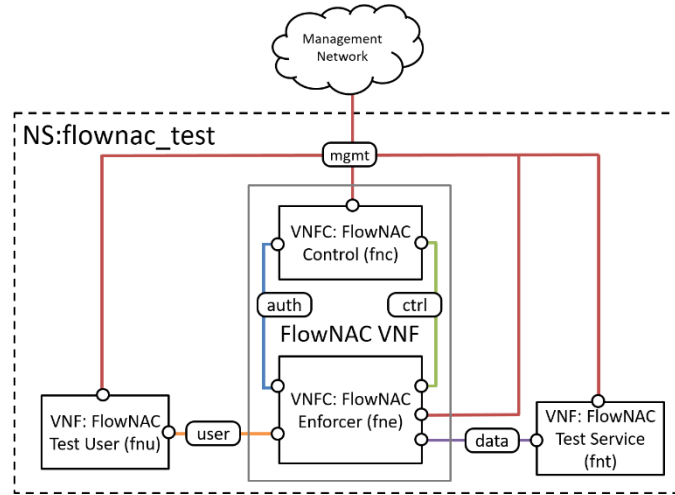


Figure 2 Testing Setup: FNC, FNE, FNU and FNT

Figure 3 shows a detailed diagram of the actual deployment for the initial setup (baseline). There is one node (vh05) for MANO and VIM, and for NFVI two compute resources (vh06 and vh07 and one networking resource that interconnects both.

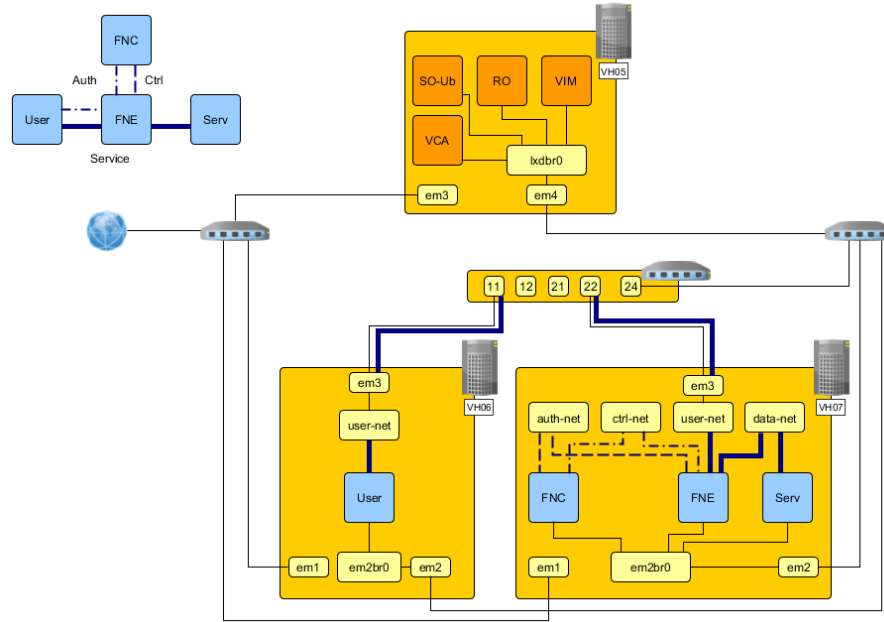


Figure 3 Detailed Testing Setup

Regarding the FNE, different alternatives will be proposed to implement its functionality:

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

There is another scenario which is planned for future instances of the PoC. Which is an evolution of Scenario 2:

- Scenario 2a – DPDK-enabled FNE VNF

Current NFV architectural framework does not define clearly the network virtualization solution and the proper interface (isolated) to interact with the underlying SDN fabric (NFVI).

A.2.3 Mapping to NFV ISG Work

GS NFV-EVE 005 (Report on SDN Usage in NFV Architectural Framework) shows the different approaches to integrate SDN in the NFV architectural framework, but most of the content is related to the fact that the underlying network can be controlled by software (by means of a SDN framework) to dynamically create links or route traffic at the infrastructure or tenant level. In this document, the PNF is proposed to implement the SDN controller but not the data path as proposed in this PoC. It's important to note that an IEEE Software Defined "SDN in NFV Architectural Framework" (<http://sdn.ieee.org/newsletter/may-2016/sdn-in-nfv-architectural-framework>) is a widely cited document that reflects this.

We think that stateless processing (as for now available on OpenFlow switches as filtering or header field modification) of information packets at the data path level over dedicated communication hardware (switches) should be handled by NFV architecture in a more clearly and idempotent to computing approach.

Moreover, we foresee that the P4 language¹, and more probably its chip incarnation in Barefoot Tofino², or other approaches like OpenState SDN³ will provide an evolution towards a stateful network packet data path processing that should be taken into account by the NFV architecture. Additionally, the deploying of communications NF like routing and switching over communications equipment (virtual switches or virtual routers) could benefit from this approach.

We consider that these are open issues that should be addressed as soon as possible to maintain the impact of ETSI NFV ISG in the industry.

Figure 4 shows how the components of this PoC are mapped into the NFV architecture.

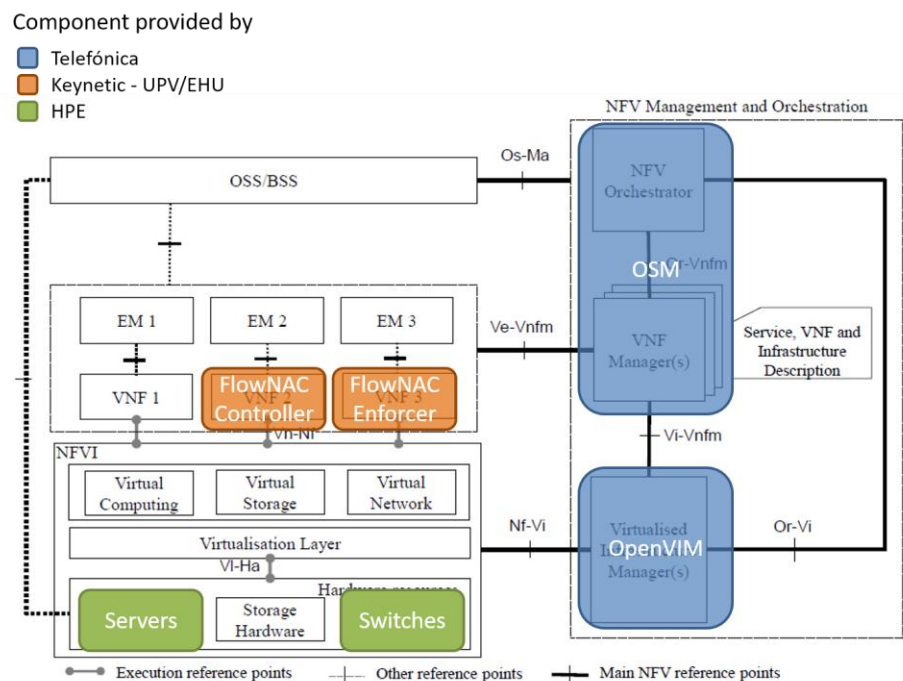


Figure 4 PoC components mapped to the NFV architecture

Describe how this PoC relates to the NFV ISG work:

- Specify below the most relevant NFV ISG end-to-end concept from the NFV Use Cases [Error! Reference source not found.], Requirements [Error! Reference source not found.], and Architectural Framework functional blocks or reference points [Error! Reference source not found.] addressed by the different PoC scenarios:

	Use Case	Requirement	E2E Arch	Comments
Scenario 1, 2 and 3			All functional blocks (NFV Orchestrator, VNF Manager, Virtualized	Illustrates interactions between the different components of the architectural framework

¹ <http://p4.org>

² <https://www.barefootnetworks.com/technology/>

³ <https://github.com/OpenState-SDN>

A.2.4 PoC Success Criteria

The PoC demonstrates in a practical use case the benefits from using the proposed criteria (statefull network function components and stateless data path processing components) to separate the VNF into sub-components.

Moreover, the different ways to implement the stateless data path processing components are shown to evaluate the pros and cons of each option, ranging from performance to overall resource usage efficiency.

A.2.5 Expected PoC Contribution

One of the intended goals of the NFV PoC activity is to support the various groups within the NFV ISG. The PoC Team is therefore expected to submit contributions relevant to the NFV ISG work items as a result of their PoC Project.

List of contributions towards specific NFV ISG WIs expected to result from the PoC Project:

- PoC Project Contribution #1: Demonstrating the decomposition of a VNF into separate sub-components: stateful network function components (compute) and stateless data path processing components (networking). NFV WI IFA029
- PoC Project Contribution #2: Demonstration of efficiency improvements in VNF data processing while minimizing the overall resources used (resource occupation) from the NFVI, mainly when considering stateless data path processing components. NFV WI IFA025
- PoC Project Contribution #3: Contribution to support VNF functionality splitting (between compute and networking) under ETSI NFV scope.. NFV WI IFA029
- PoC Project Contribution #4: Feed results related to the implementation of the orchestration of the decomposed VNFs and VNFCs to open-source projects, and specifically OSM. NFV WI TST003