

# Silicom NFV Solutions

---

*Virtual Service Chaining, Close to the Edge*

---



## Network Functions Virtualization

Much was discussed and described with regards to NFV, especially with regards to NFV software platforms and NFV switching and network management; and not without good reason. For one thing is true for sure (before delving into what NFV is all about) - It is a lot about software. However, designing standard hardware platform that would support the requirement of NFV software systems can be done in accordance to basic understanding of NFV building blocks.

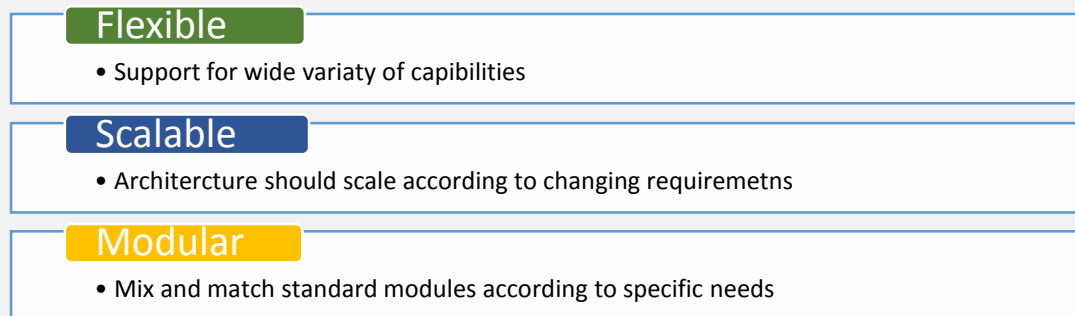
Understanding the big picture of NFV requirements, can be achieved by analyzing what NFV is comprised of. This analysis divides the NFV framework into three major prominent

characterizations, where for each characterization further break down is done to complete the details of the big picture. According to this analysis Silicom devised its offering for NFV.

## Understanding NFV

### NFV Mission Statement

NFV as a term encompasses diversified technologies, some of which are open, while others are commercial and less open. However, all this collection of technologies had and still has a clear target – **enabling service providers to provide better service for their end customers, while maintaining competitive edge**. This mission statement immediately sets clear requirements. This platform to be used by service providers, should thus be:



To address those requirements, three high level features common to all NFV deployments gradually took shape and are now the elements that identify and describe NFV deployments. NFV deployment in deed involves **service virtualization**, mainly for **service chaining**, while it would all reside **close to the network edge**.

### Virtualization

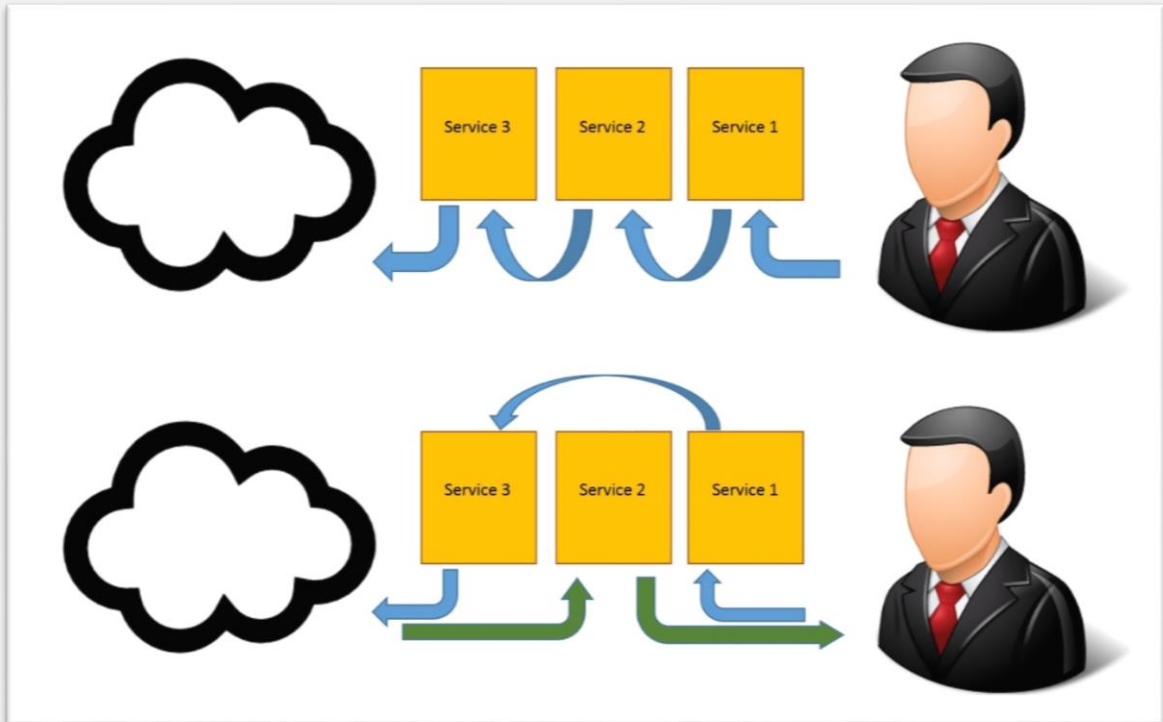

The building block of the service that is provided by a provider to a customer may be embedded on two distinct layer. At the top, we will find network applications layer such as security applications, monitoring applications or other; and right underneath there would lie the platform layer, where platforms are the service that is provided to the customer. Rendering the above as a service for customers create the software as a service



(SaaS) and platform as a service (PaaS) paradigms; instead of application and platform, respectively, delivered on appliance. The idea is to have the same functionality of a network application, adapted to the requirements of the end customers, running and managed on behalf of the end customer, on resources that are managed by the provider. Hence, virtual network function (VNF) came about as a generic term describing a piece of capability (of a network application or of a software platform) that is invoked to serve the end customer. Those VNFs that originally were born and raised on their own hardware platform, are groomed to accommodate in a flat standard environment, one fits for all.

### Service Chaining

In order to expand business, service providers wish to expand capabilities of addressing end customer needs, with the widest possible features set, that would occupy the lesser extent of real estate. End users diversify in requirements therefore each SLA may differ from the next. In order to address exactly that, service chaining is adopted, as a means to deliver the right service to the right end customer, according to the right SLA.



*Figure 1- Types of Service Chaining*

Services may vary, and may contain security services, cache services, content delivery, etc.

### Close to the Edge

The purpose of serving end customer with the closet proximity to end customer network is to deliver the best service. CDN or content cache services, for instance, gains their advantage out of being close to their consumer. As compute and networking power are getting more commoditized, and more network applications are reincarnating themselves as virtual instances, being close to the edge becomes possible, and the benefits the stem out of it can become of use.



## Silicom Network Appliance NA2600

Compute and networking integrated in one. A COTS server mother board of choice, combined with switching fabric, Silicom's Network Appliance deliver a tailor made yet standard platform, to accommodate virtual network functions.

Flexible - Off the shelf mother board of choice

Scalable - Scales well both with VMs and cascaded server

Modular - Front connectivity mix and match

Network Appliance NA2600 is offered with **flexible** choice of standard, commercial off the shelf x86 based motherboards, with single or dual socket CPUs. Its front connectivity can be defined as a choice and may include a combination of 10GbE, 40GbE or 100GbE. Further designs can be **modularized** according to specific requirements.



*Figure 2 - Network Appliance NA2600*

### Use Cases

#### Cellular Backhaul and GTP Tunneling

The challenge of GTP processing is being able to consistently identify GTP-U tunnel by TEID. Since all GTP sessions between GGSN and SGSN [1] are carried over same IP address and UDP port, then five tuple filtering is not enough to identify user session within tunnel – but it is a start. Supporting redirection of GTP traffic according to TEID can be achieved with NA2600 in three configuration steps:

- 1) 5 tuple based filtering and redirecting by the 5 tuple engine of the switching fabric, of GTP-C traffic to host's compute
- 2) Host identify the TEID of the opened tunnel and set the TCAM engine the switching fabric accordingly to filter and redirect TEID based GTP-U traffic
- 3) Further GTP-U traffic is processed, filtered and redirected on host's compute

Once GTP-U tunnel is TEID'd it then can be forwarded and chained into configured services.

## Service Chaining

As stated above, service chaining is one of the three fundamental aspects of NFV solution. A scalable forwarding mechanism for purposes of service chaining can hardly rely on forwarding or switching overlay. Each end user, of which traffic should be chained across different services, requires a dedicated tunneling mechanism; and this is cumbersome to achieve by means of L2 tunneling (VLAN, VxLAN, etc.) or L2/L3 encapsulations such as NVGRE, due to the fact that one end customer's chaining could differ from the next's.

There are attempts to address this exact issue and even to standardize it [2]. However, similar approach of this standardization can be found in NA2600 as an infrastructure. On top of both standard and advanced forwarding mechanisms that are implemented in the switching fabric of NA2600, there is the inner forwarding mechanism, unique to the fabric. The mechanism if controlled by specific headers that are prepended to the traffic, and stripped on the ingress. Taking advantage of this generic tunneling capability, a scalable high performance service chaining mechanism can be implemented.

Once identified and classified by host, customer traffic can be assigned a proper header that would lead this traffic's forwarding within the NA2600 according to the desired chaining. When forwarded onward from the system, this header would be stripped off the packet by the fabric itself.

This way, chaining would be performed by a standard yet dedicated hardware that is built for that purpose, relieving CPU to perform its compute tasks.

---

## Resources

[1] [www.etsi.org](http://www.etsi.org)

[2] <https://tools.ietf.org/html/draft-niu-sfc-mechanism-01>