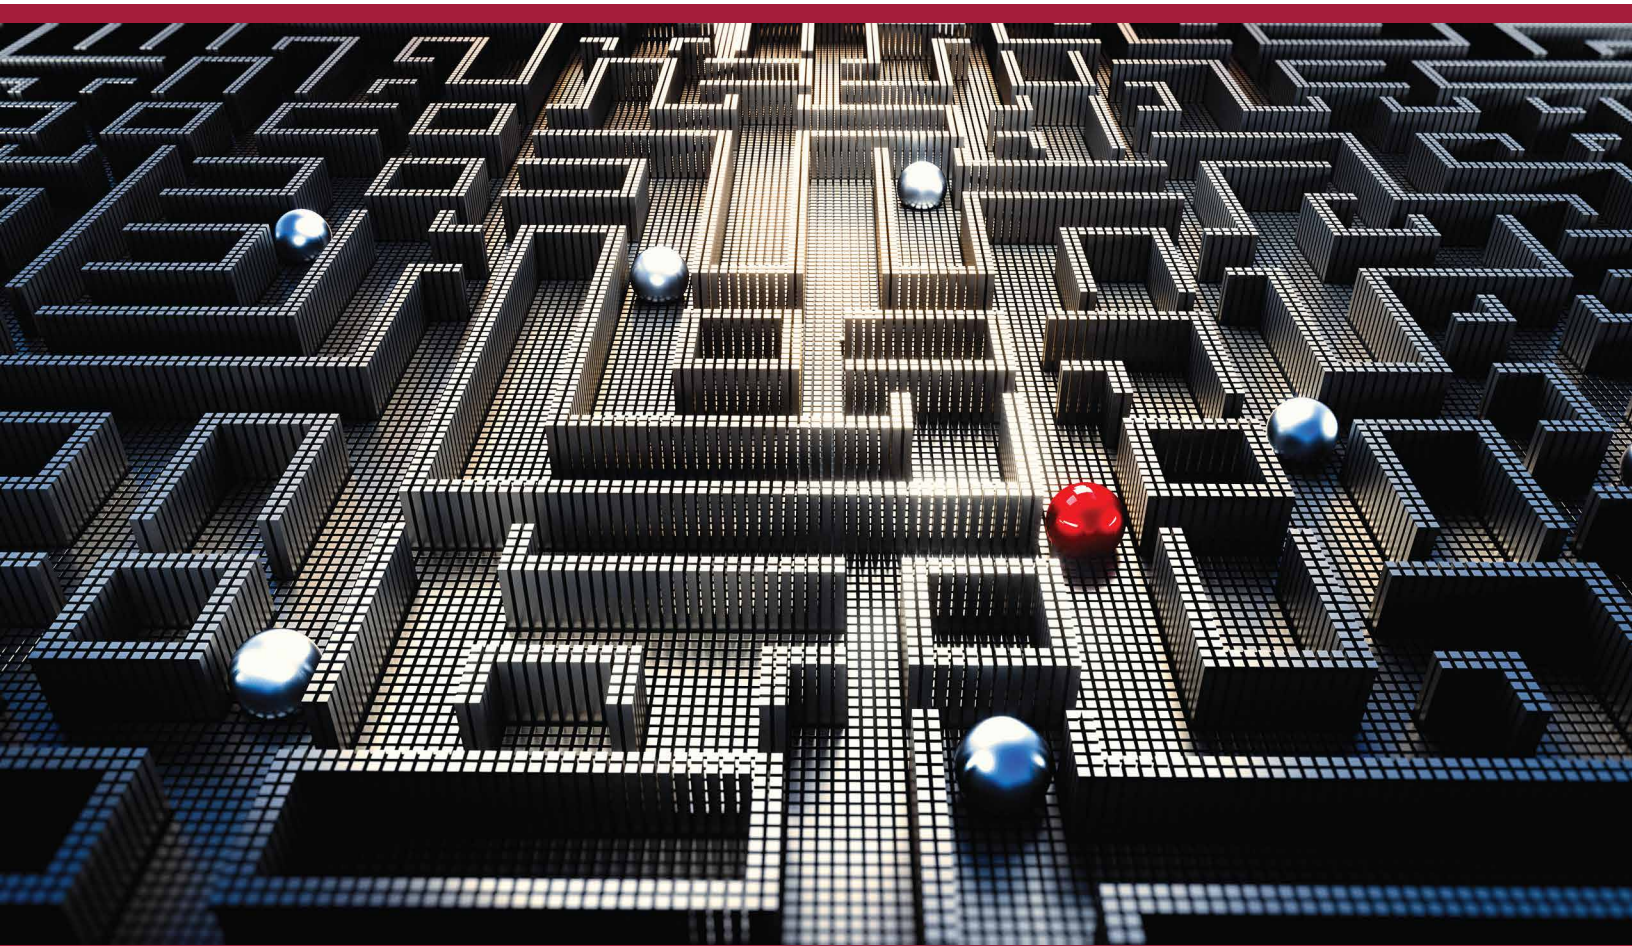




THE READY FOR  
ANYTHING NETWORK

WHITE PAPER

# Lessons Learned on the NFV Front Lines



# Lessons Learned on the NFV Front Lines

## Contents

- Lessons Learned ..... 3
- Lesson #1: Not all hardware is created equal ..... 3
- Lesson #2: The packet forwarding architecture and hypervisor need attention too ..... 3
- Lesson #3: Don't oversubscribe the application ..... 4
- Lesson #4: NFV isn't a simple plug-and-play solution ..... 4
- Lesson #5: Redundancy needs to be built into the application, not just the NFVI architecture .... 4
- Lesson #6: Telecom applications require built-in load balancing ..... 5
- Lesson #7: VMs need to scale independently ..... 5
- Lesson #8: Ownership is important ..... 6
- Lesson #9: One EMS is better than two (or three) ..... 6
- Lesson #10: Learn from the leaders ..... 6

[www.affirmednetworks.com](http://www.affirmednetworks.com)

© 2018 Affirmed Networks. All rights reserved




## Lessons Learned

It is the opinion of some telecommunications network equipment vendors that Network Functions Virtualization (NFV) is a byproduct of 5G, and that the one shouldn't arrive before the other. Reality says otherwise; many communication service providers are deriving value from NFV initiatives right now, primarily in the form of capex/opex savings and network agility. Yet many service providers, in our experience, still only tap into 30 to 40 percent of NFV's true potential.


What prevents providers from realizing NFV's full potential? There is no single reason. Rather, it's a combination of missed opportunities and misunderstanding as to NFV's architectural requirements. In the last eight years, Affirmed Networks has helped leading service providers successfully transition to NFV based architectures and realize exceptional returns. Along the way, we've learned some valuable lessons on how providers can avoid underwhelming NFV results and realize the technology's full transformative benefits.

### Lesson #1: Not all hardware is created equal



The belief that you can run virtualized telecom applications on any vendor's server is only a half-truth. There is one hardware dependency that always needs to be considered: the hardware must have a network interface card (NIC) that supports the data plane development kit (DPDK) in order to function properly. In our experience, we've found it's often better to bundle the virtual network function (VNF) with hardware providers that support this NIC requirement rather than deploy the VNF in a hardware-agnostic environment.

### Lesson #2: The packet forwarding architecture and hypervisor need attention too



While choosing the appropriate hardware can aid in the performance of your virtualized network, the packet forwarding architecture requires attention as well. The main function of the evolved packet core (EPC) is to move a large number of packets through the data plane. This means you need very high performance in the data plane. Typically, packets travel through the vSwitch function within the hypervisor, which queues them for the virtual machines (VMs). The vSwitch function uses a great deal of computing power, which limits the performance that VMs can achieve. This creates a need for single-root input/output virtualization (SR-IOV) technology to get around this limitation. SR-IOV technology allows the packets to bypass the hypervisor layer and travel directly from the PCI on the server to the VMs, giving the VMs full use of all CPU power and significantly increasing performance. While SR-IOV is not a requirement for NFV deployments, its role and impact are sometimes misunderstood by



service providers. If a provider requires very high throughput, then SR-IOV is necessary. Furthermore, applications are very sensitive to how the hypervisor is configured and the specific settings it uses. In order to reach maximum performance, service providers must also tune the hypervisor to meet the specific requirements of their application (e.g., tuning how the hypervisor schedules the CPUs, CPU pinning, etc.).

### **Lesson #3: Don't oversubscribe the application**



Another important lesson learned is never oversubscribe a virtual application or the application's CPU. Even though the technology allows for oversubscription of the application, this ends up degrading the performance of the application and causes problems down the road.

### **Lesson #4: NFV isn't a simple plug-and-play solution**




Virtualization is often marketed as plug and play, but in reality it requires some tuning in the ecosystem for telecom applications to run at maximum performance. For example, in one customer deployment, they experienced a denial-of-service attack that featured a lot of "burstiness" in the traffic. The DPDK driver was indiscriminately dropping packets and causing packet loss because it didn't have any concept of quality of service (QoS). This required modification of the driver to avoid latency and packet loss. While this may seem like a minor detail, it can have a major impact on the performance.

### **Lesson #5: Redundancy needs to be built into the application and not just the NFVI architecture**



In the enterprise world, redundancy is a relatively simple matter of spinning up a new VM when one VM fails. This works well for stateless, transaction-based applications, but telecom applications are stateful. When you lose the state of the VM, you lose the service. Also, when a VM fails, the time it requires to spin up a new VM is far too long for telecommunications applications and extends the problem of service disruption. In order to provide stateful redundancy in a telecom environment, operators cannot rely only on NFVI redundancy; statefulness needs to be built directly into the virtual application itself or maintained in an externalized database. That's the approach we took when building our virtualized EPC solution, and it is a very important lesson to remember when talking about NFV.

## Lesson #6: Telecom applications require built-in load balancing



One of the main benefits of a virtual environment is the ability to scale up or scale down your processing power as workload demands change. When decommissioning a VM, however, you lose the state of that VM. In an enterprise environment featuring stateless, transaction-based applications, this is not an issue—but it is an issue in a telecom environment where stateful applications are the norm.

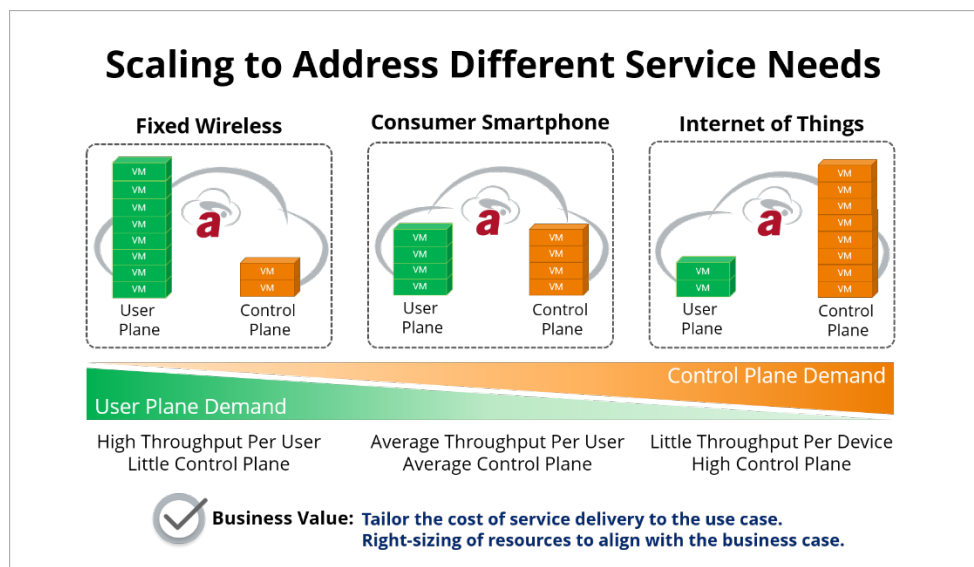
Telecom applications that support dynamic scaling need load balancing; this way, when new resources are available, the application can load-balance across the new resources to prevent dropping service during a call/session. We believe load balancing should be built into the application, as the application knows better how to use the resources than an external load balancer.

## Lesson #7: VMs need to scale independently



Scalability is something NFV vendors need to be thinking about before they build their solutions, not after. Specifically, vendors need to ensure that their VNFs can scale independently across different dimensions. In a telecom application, the data plane, management plane and control (i.e., signaling) plane each need to be scaled independently to avoid paying for stranded capacity. In a blade-based architecture, the signaling, data and management capacity are added in fixed ratios; as more signaling capacity is needed, more blades are added. The result is that service providers end up with more data capacity as well, whether they need it or not.

In a virtualized architecture, where independent scaling is supported, providers can scale up signaling capacity without affecting the data or management dimensions. This is why we chose to decompose each plane when we built our vEPC.



Applications need to be designed in a flexible way, allowing the scaling of VMs based on the specific call model or application (e.g., IoT, enterprise, consumer) and the availability of resources. By doing this, service providers can right-size the capacity for the specific call model.

## Lesson #8: Ownership is important



Traditionally, service providers have relied upon their vendors to provide all the layers of a solution. An NFV architecture is different. There's a hardware layer, a hypervisor layer and an application layer to consider, with each vendor bringing their own perspective to the solution. Instead of one finger to point when things go wrong, providers must now point several fingers. This creates a challenge for service providers in managing deployments, as there is no clear accountability. At Affirmed, we've countered this problem by taking "ownership" of the NFV experience and ultimate responsibility for the way our vEPC solution behaves in the NFV infrastructure (NFVI) environment. Our customers appreciate having an experienced vendor as a lead implementor who can work with ecosystem partners to resolve any issues.

## Lesson #9: One EMS is better than two (or three)



Service providers are accustomed to a single element management system (EMS) that displays the state of the system (e.g., alarms, traps, etc.) across all solution layers. In an NFV architecture, however, there are separate element managers for each layer. Having an overarching EMS that extends visibility into all layers and manages them as a single pane of glass" is an important capability for any NFV architecture.

## Lesson #10: Learn from the leaders



Perhaps the most important lesson there is to be learned from the leaders in the NFV journey is not to wait. There are those vendors who will tell you that NFV isn't ready for prime time. What they're really saying is that their solutions aren't ready yet. At Affirmed, we're building virtualized solutions that give the leading operators of today the competitive advantage they need to remain the leaders of tomorrow. Our cloud-native, 5G-webscale solution not only reduce capex and opex, but also provide the capabilities for new revenue-generating services including service automation and microservices creation. For more information on how Affirmed Networks can help you move ahead, visit us at [affirmednetworks.com](http://affirmednetworks.com).