THE READY FOR ANYTHING NETWORK

WHITE PAPER

Economic Impact of NFVI Design Choices

Learn the Four Critical NFVI Design Choices that will Impact the Cost and Performance of Your Network

Economic Impact of NFVI Design Choices

Learn the Four Critical NFVI Design Choices that will Impact the Cost and Performance of Your Network

Contents

www.affirmednetworks.com

Economic Impact of NFVI Design Choices

"It's important to understand that certain NFVI design decisions can have a material economic impact on the overall solution cost." – Tim Irwin, VP of Pre-Sales Systems Engineering @ Affirmed Networks

Telecommunications carriers are moving to software-defined networking (SDN) and network functions virtualization (NFV) architectures as a means to deliver new services more efficiently, scale their networks more effectively and drive down costs. Yet, in focusing on NFV/SDN as the goal, carriers may risk losing sight of the fact that how they arrive there—that is, the network design choices they make along the way—can have a significant impact on network cost and performance. In fact, just making one or two seemingly minor changes in the design of the NFV infrastructure (NFVI) can end up saving (or costing) carriers millions of dollars down the road.

For these reasons, experienced guidance is critical when designing an NFVI architecture. Unfortunately, many carriers do not have a lot of experience in building SDN/NFV environments; a problem that is compounded when carriers bring together IT and telco application teams to create an SDN/NFV design plan. Because of the very different network requirements of IT and telco applications, this approach can quickly create confusion instead of synergy. Bringing in outside expertise is certainly an option, but here again carriers need to be careful. Many vendors have an ulterior motive for advancing their own NFV/SDN approach, particularly when hardware sales are at stake, and carriers may miss an opportunity to realize the full economic benefits of NFV/SDN as a result.

Critical Design Decisions That Matter

There are four critical areas in the functional design of an NFVI architecture that can have a profound impact on cost and performance:

1. Packet-forwarding method
2. Data center design
3. Software efficiency
4. Hardware choices

It's not uncommon for carriers to identify and address one or two of these areas in their own design efforts but, in our experience, it's very rare to encounter a carrier that has taken all these areas into account when designing their NFV network. This is unfortunate, because making a less-than-ideal design decision in any of these areas will have a significant impact on network and services costs down the road.

Design Decision #1: Packet-Forwarding Method

On the surface, combining IT and telecommunications groups and tasking them with the design of an SDN/NFV architecture might seem like a good idea. Packet forwarding is a prime example of where this approach can be problematic, as IT applications behave very different than telco applications in several important ways:

1. Telco applications emphasize packet-forwarding performance because the packet core network is basically a middle man: it receives the packet on one side and sends it out the other side. IT applications are point to point, so packet forwarding is less important.
2. IT applications are stateless; if the connection fails, they simply re-send the packet. Telco applications are stateful; if an E911 call fails, it's a big deal. Similarly, it's possible to "overbook" IT applications in a virtualized environment by combining a lot of lightly used web servers onto the same physical server and moving an application to a new server if it gets busy. You wouldn't be able to pre-empt a voice call and spin up a new server to host it.
3. IT applications are typically CPU intensive rather than I/O intensive, so they'll often max out the CPU and have I/O capacity left over. Telco applications behave in the opposite manner. If you size your NFVI network I/O capacity based on the CPU capacity you have—a common tactic in the IT world—you can end up with a considerable (and costly) amount of stranded CPU resources.

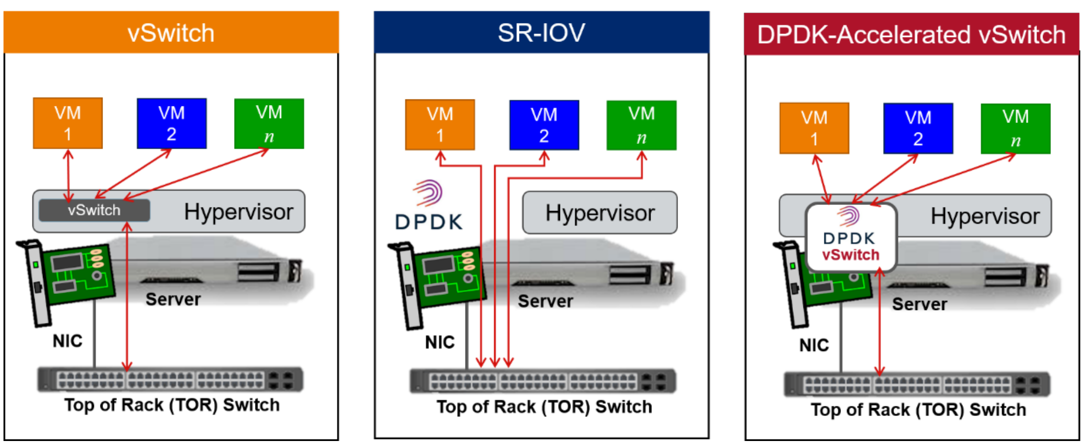

Figure 1. Classes of NFV Packet Forwarding Solutions

There are three different ways to approach packet-forwarding in an NFVI architecture: vSwitch, Single Root Input-Output Virtualization (SR-IOV) and Data Plane Development Kit (DPDK) Accelerated vSwitch. Each method has its own set of pros and cons, as explained below.

- An **integrated vSwitch** approach is commonly used for IT applications. It supports SDN functionality and allows for virtual machine (VM) mobility but delivers substantially lower performance compared to the SR-IOV and DPDK-accelerated vSwitch methods.

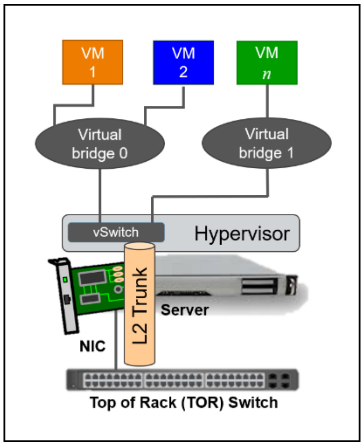

Implementations

- Linux Open vSwitch (OVS)
    - Open-source Linux implementation
    - Default for OpenStack
- VMware vSwitch
    - Default for VMware deployments
- Very mature

Emphasizes L2 bridging

- VLAN 'Q-in-Q' approach
- But may also be used with OpenFlow

Data Center Networking

- May be used with traditional L2 design (spanning tree)
- May be used with SDN overlay (VXLan, Contrail, etc.)
- Good support for VM mobility

May incorporate L3 host NAT

- Not appropriate for packet forwarding NFV use

Packet Forwarding Performance

- Poor due to Hypervisor kernel interrupts

Figure 2: Integrated vSwitch

NFV network engineers will be more familiar with the SR-IOV approach. It delivers better performance for user-plane packet forwarding, but bypasses the hypervisor to achieve this performance, which means an SDN overlay is not possible.

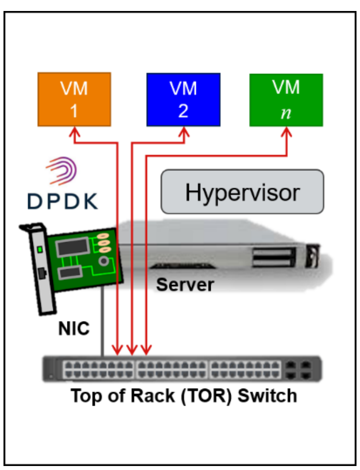

Broad ecosystem support

- Supported in both Linux and VMware hypervisors
- Very mature

Commonly deployed with Intel DPDK

- DPDK implements the Ethernet controller & packet memory buffers in User Space (as opposed to Kernel Space)
- Avoids expensive packet copy operations
- Leverages poll-mode drivers rather than interrupts

Data center networking implications

- VM VLAN membership is handled in TOR
- Can be used in conjunction with OVS for non-I/O intensive networking
- VM traffic bypasses hypervisor, therefore functions such as vTap cannot be supported
- Poor VM mobility

Packet forwarding performance

- Best performance relative to other options

Figure 3. Single Root Input Output Virtualization (SR-IOV)

The DPDK-accelerated vSwitch approach represents a compromise of the two preceding examples. It delivers better performance than a standard integrated vSwitch approach, but lower performance than the SR-IOV approach. In exchange, carriers gain the ability to implement an SDN overlay with the DPDK approach.

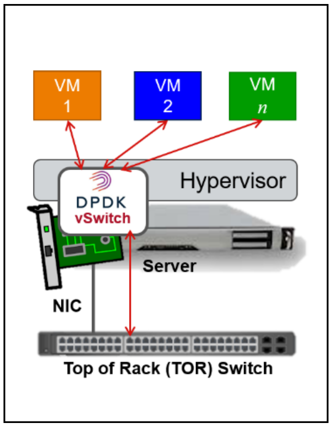

Implementations
- Redhat DPDK-Accelerated vSwitch
- Juniper Contrail vRouter
- Maturing, but improvement still needed

Data center networking implications
- Gain SDN overlay benefits (VM mobility)
- Possible to have a single, consistent design

Packet forwarding performance
- Better than vSwitch
- Worse than SR-IOV + DPDK

Performance Trade-offs
- Requires reserving some number of vCPUs for User space part of vSwitch/vRouter
- Less vCPUs can be allocated to application VMs
- Flow learning rate could be a bottleneck
    - vSwitch → SDN Controller flow classification

Figure 4. DPDK Accelerated vSwitch

A Closer Look at Performance

Packet-forwarding performance has a direct impact on network cost. As Figure 5 illustrates, I/O performance varies greatly between the different packet-forwarding methods, from 7 Gbps* on the low end (vSwitch) to 150 Gbps on the high end (SR-IOV), with the DPDK-accelerated vSwitch somewhere in the middle at 112 Gbps. (*Based on a 650 byte packet running on a Skylake series Intel server.) In our example, we consider a 150 Gbps I/O deployment. Where an SR-IOV packet forwarding would require just 1 server, the same network using a standard vSwitch architecture would require 21 servers.

Read the Affirmed/Intel Mobile Core Performance Report

Baseline VM I/O Performance Comparison (per server)

Intel Proc Architecture	Packet Size	VM Forwarding Method	Throughput (Gbps/server*)	Notes
Skylake	650 bytes	SR-IOV	150	Network I/O efficient, not processor bound
Skylake	650 bytes	DPDK Accelerated vSwitch	112	Limited availability/maturity, but improving
Skylake	650 bytes	Traditional vSwitch	7	CPU interrupt limited

* Affirmed delivers 150 Gbps throughput on single Intel Xeon Server

Example – 150 Gbps I/O Deployment (redundancy included)

VM Forwarding Method	# of Servers
SR-IOV	1
DPDK Accelerated vSwitch	2
vSwitch	21

VM forwarding architecture directly impacts overall cost and scaling
As much as a 20x difference

Figure 5. Implications of VM Forwarding Model

For many carriers, the packet-forwarding decision comes down to whether or not they want an SDN overlay in their network. This raises the question: Why does having an SDN overlay impact performance? The answer is twofold. In an SDN, you're handling the packet multiple times, which requires additional processing and adds latency. Also, between OpenStack's agents and DPDK's vSwitch/vRouter requirements, you're consuming about 25% of your server capacity even before you load a single VM onto it.

It's important to note, however, that SR-IOV isn't always the "best" choice simply because it's the fastest. Carriers need to weigh the pros and cons of each option and, ideally, optimize their NFVI architecture around different workload types. We certainly don't advocate creating a dozen different NFVI architectures in your network, but we also believe that there are beneficial tradeoffs to a heterogeneous NFVI approach versus a single, homogeneous approach.

The effect of different NFVI architectural designs can seem minor when viewed through the lens of a low-bandwidth IoT application. Those differences will appear much more significant, however, as you look at high-bandwidth 5G services that may consume as much as 20 Gbps per subscriber in the user plane. In such a case, there would be a very compelling argument for using DPDK or even standard vSwitch packet forwarding for IoT applications, and SR-IOV packet forwarding for applications in the user plane.

Design Decision #2: Data Center Design

For telco applications, everything centers around the evolved packet core (EPC) architecture. The EPC is essentially the heart of telco services in the data center. Within a virtualized EPC (vEPC), you'll encounter three different types of VMs: user plane VMs, control plane VMs and management VMs. Current NFV/SDN guidelines recommend implementing control and user plane separation (CUPS) to improve network performance and agility. Many network vendors are looking at ways to support CUPS in their current solution set. The Affirmed vEPC, by contrast, has supported the separation of control and user planes from day one.

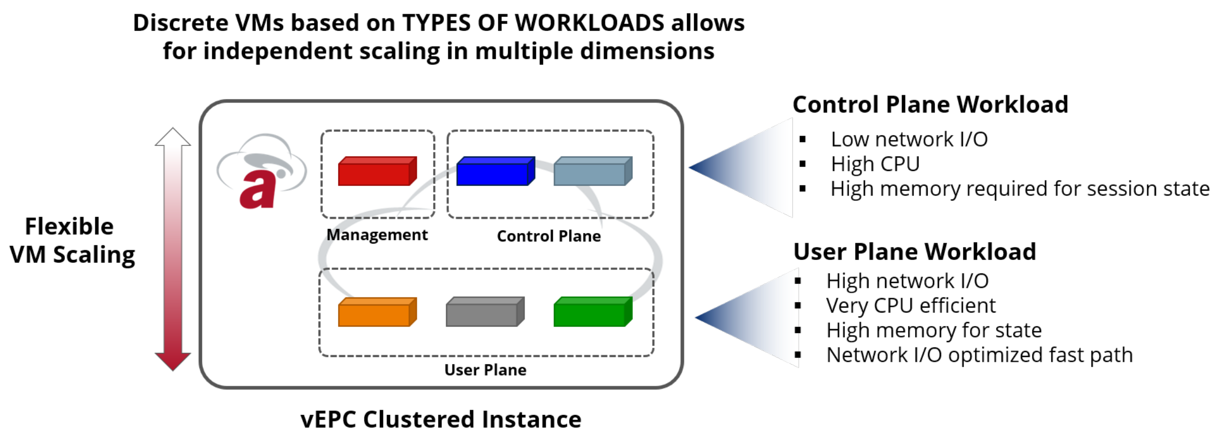

Figure 6. EPC Workloads

Why is control and user plane separation important? Because these plane workloads tend to behave very differently. Control plane workloads mimic traditional IT applications in that they tend to consume a lot of CPU and not so much I/O throughput. In the user plane, the opposite is true. Also, there are opportunities to improve latency by moving user plane functions closer to the user. In a traditional EPC architecture, the control and user plane functions are grouped into a single, centralized cluster.

With CUPS, carriers can now split these functions into separate clusters, so that control plane functions are centralized in one cluster, and user plane functions can be distributed closer to the network edge in a separate cluster. This approach also allows carriers to create different NFVI architectures tailored to the workload characteristics of each cluster, including optimizing servers for the unique requirements of the workload, which can reduce capex costs by up to 60%.

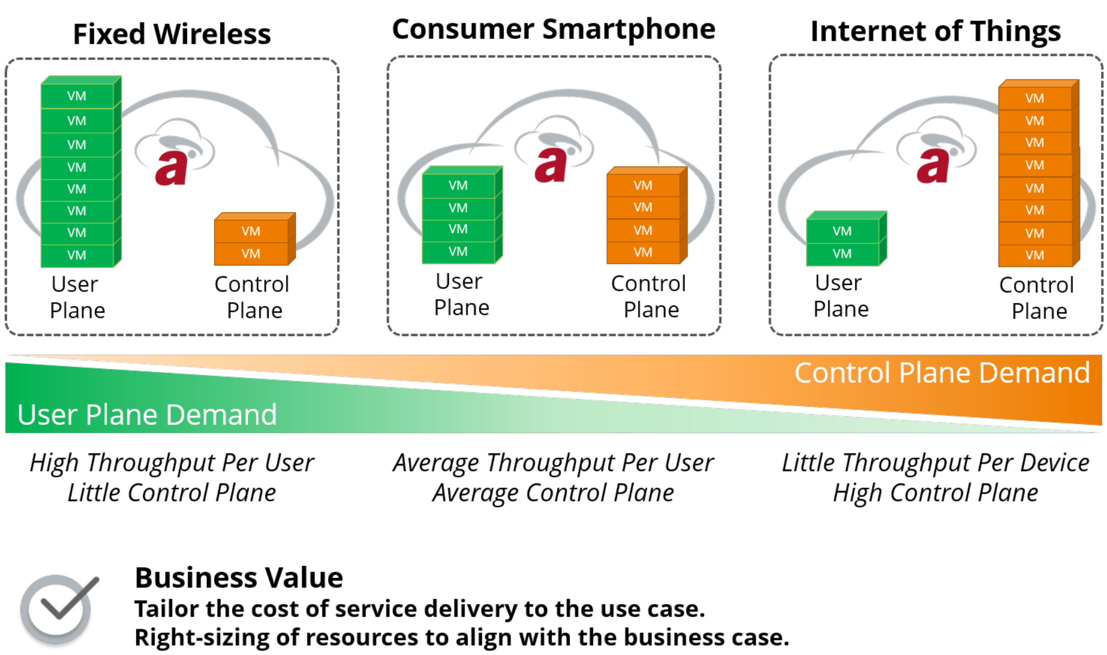

Figure 7. Scaling to Address Different Service Needs

Design Decision #3: Software Efficiency

Integrated Vs Standalone Network Functions

Historically, carriers have taken a “mix and match” approach by choosing discrete products from different vendors such as signaling gateways, packet gateways, voice application servers, firewalls, etc., and putting them under the direction of different teams—i.e., packet core managed by one group, and GiLAN VAS by another. While this does enable operators to choose different vendors as they build their network, the result is that packets must be passed up through the application chain multiple times. By repeating the same steps each time, the process consumes a considerable amount of CPU and adds latency for each additional hop through the chain. In some cases, as much as sixty percent of the packet processing time can be attributed simply to moving the packet from the Network Interface Card (NIC) through the different elements.

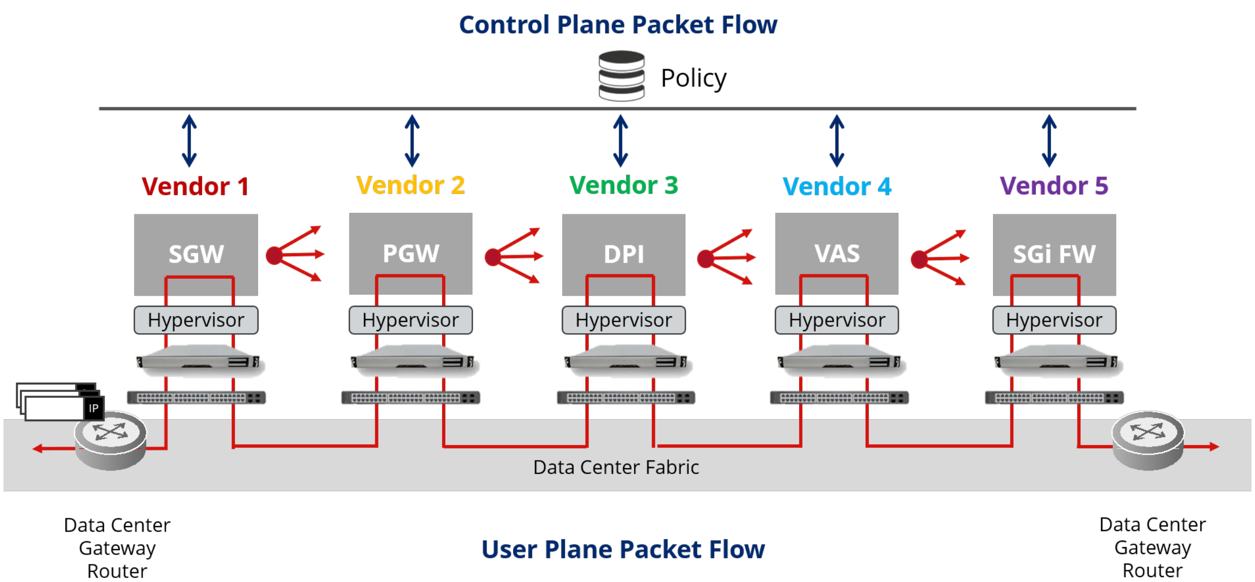

Cons

- *Increased control plane impact on scaling.*
- *Managing capacity differences between functions can be a challenge.*
- *Possible overlap between functions (e.g., standalone DPI + FW DPI).*
- *Incurs additional overhead & latency in the user plane since every packet must be handled multiple times.*

Figure 8. Discrete Functions: Cons

With an integrated NFV solution, network functions are combined so that the packet only needs to be brought up in application once and all of the processing is done at the same time. The analogy we can use is that of taking multiple bites of the apple at once rather than separate bites each time. This approach reduces both the processing demands and the latency of the packet (see Figure 9). The downside of a combined function approach is that only a few vendors currently have the capabilities to support this.

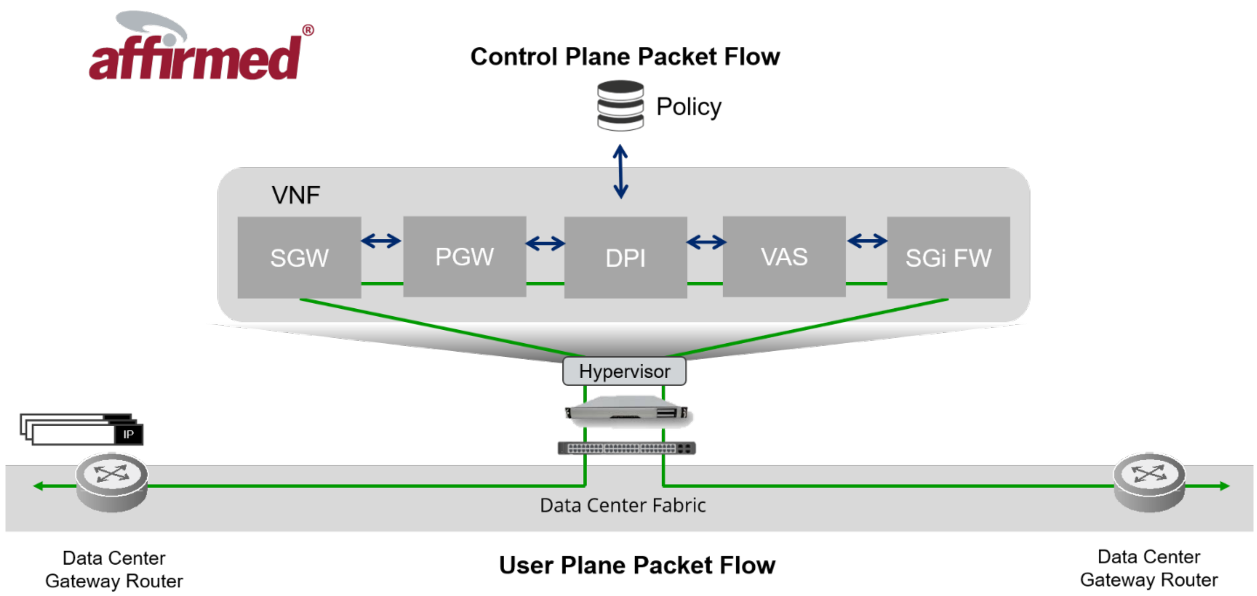

Pros

- *Reduction in external signaling – information queried once and shared among functions.*
- *User plane packet is pulled into memory once.*
- *All necessary operations take place at once.*
- *Decreased latency compared to multiple discrete functions.*

Figure 9. Combined Functions: Pros

Virtual Probes

As the industry has moved to NFV, most network equipment vendors have launched *virtualized* software-based versions of their hardware-based products. All virtualized software is not created equal, however; some are architected specifically to leverage the efficiencies of an NFV architecture, while others are merely software adaptations of legacy products. This distinction is clear in the case of network probing software.

Traditionally, network probe solutions have been physically connected to the device interfaces they probed with an optical tap. When you move to a virtualized architecture, it's difficult to replicate this approach. What many vendors have done instead is to create a kind of passive virtualization where a virtual tap (or vTap) now sits on the hypervisor interface. This vTap copies all of the packets sent by the VM and forwards those copies to a separate processor for storage and analysis. In effect, what this does is doubles the amount of I/O traffic while cutting server performance almost in half. And this approach is completely ineffective in an SR-IOV implementation because the hypervisor is bypassed altogether in SR-IOV packet forwarding.

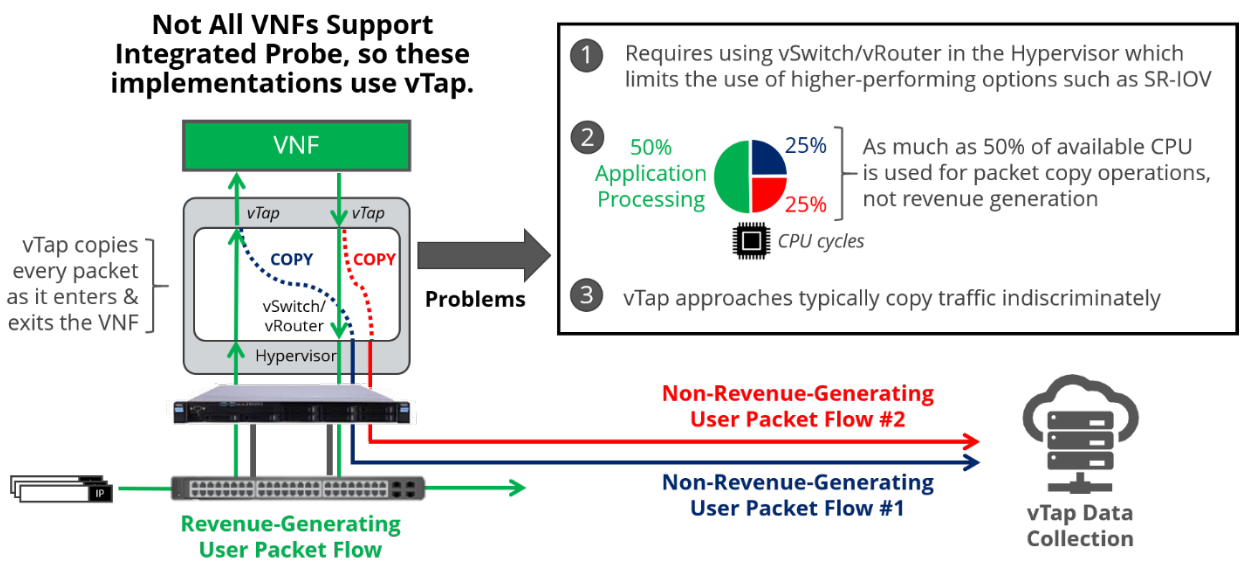

Figure 10. Virtual Tap (vTap)

Affirmed Networks took a completely different approach, creating a separate virtual network function (VNF) for network probing, which we call vProbe. Instead of mirroring data, which is highly inefficient, the vProbe VNF forwards the pre-correlated metadata to the carrier's data lake or other data repository. You're not doubling the amount of I/O traffic, you're not cutting server performance in half. In fact, activating the vProbe function has very little impact (+/- five percent) on a server's performance.

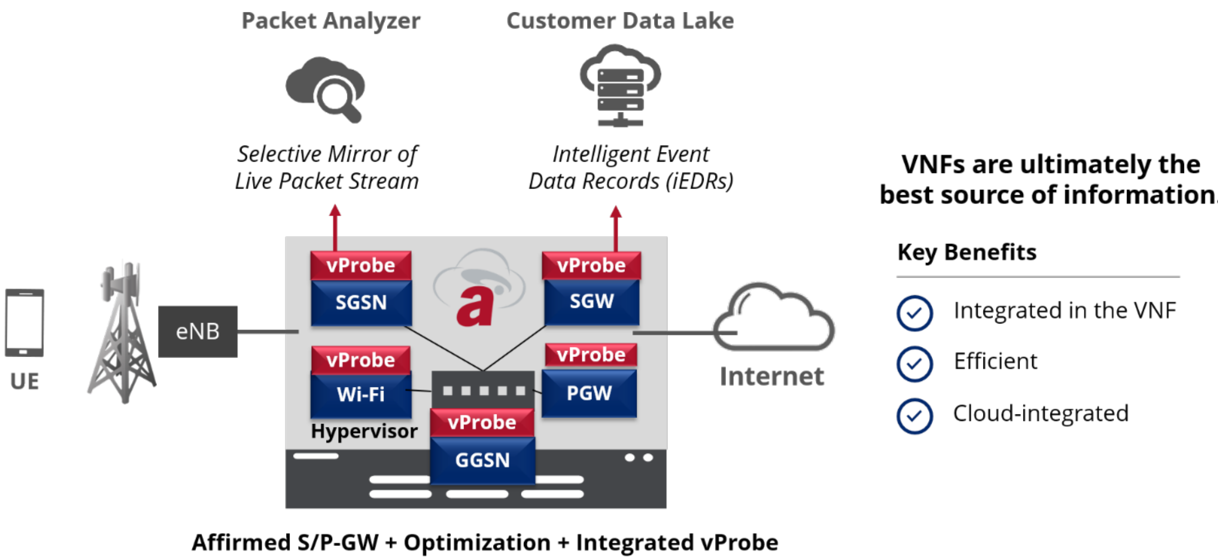

Figure 11. Integrated vProbe

Design Decision #4: Hardware Choices

While the idea of NFV is to commoditize the hardware, that doesn't mean that hardware decisions are unimportant in an NFVI architecture. Using the same server configuration for different workloads, or using an inefficient configuration, can substantially increase the number of servers required for the job—by as much as 300%. For example, let's consider the hardware for our user plane cluster. Let's say we have 100 servers in that cluster and we need to double the amount of I/O capacity. Would it make more sense to buy 100 more servers or to upgrade the 10 GbE NICs on our existing servers to 25 GbE NICs? Clearly, upgrading our existing servers provides a much more cost-effective solution for solving our I/O capacity problem.

#NIC Ports Per User Plane Server (NIC level redundancy)	4 x 10G	8 x 10G	8x25G
VM Forwarding Method	SR-IOV	SR-IOV	SR-IOV
N/S Throughput / Server	10	20	100
Total # of Servers Required to Achieve 100 Gbps (With Port Level Redundancy)	10	5	1
Conclusions	I/O Constrained, CPU Inefficient	I/O Constrained, CPU Inefficient	I/O Efficient, CPU Efficient
Example HP DL380 	**10x**	**5x**	**1x**

Figure 12. Implications of NIC Choice

Conclusion

Carriers must recognize that the NFVI design choices they make today will have an important impact on the cost and performance of their network tomorrow. And this impact is not relegated to bottom-line growth because it can affect top-line revenue too. If you're paying too much to deliver 5G services, your customers will also end up paying too much for them, and that will undermine your ability to compete.

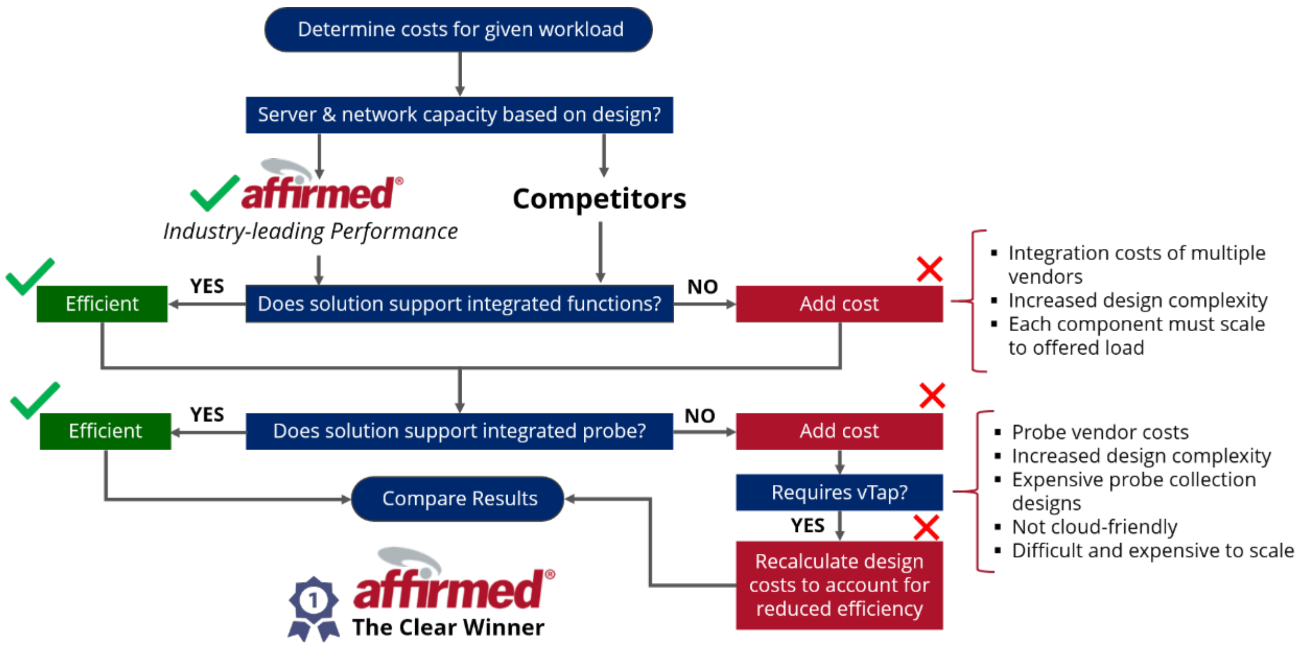

Figure 13. Economic Impact of NFVI Design

It's important to remember that no single design approach is right or wrong in and of itself. In fact, the ideal NFVI design is one that adopts different design approaches based on specific workloads rather than a single, homogenous NFVI architecture. For that reason, the best strategy for carriers is to engage NFVI experts early in the process. Affirmed is uniquely qualified to help carriers design an NFV/SDN architecture that gives them the best price/performance and aligns with the cloud/NFV standards of tomorrow. Many of the foundational architectural principles that we developed as NFV pioneers—e.g., combining virtual network functions, decoupling the user and control planes, integrated virtual probes—are today recognized by the industry as best practices for an NFV architecture. Affirmed also actively partners with the industry's NFV/SDN leaders (VMware, RedHat, WindRiver, Juniper, Intel, HPE and Dell) to deliver integrated solutions that build upon existing technology and accepted standards.

To learn more about Affirmed Networks' cloud-native NFV/5G solutions and services, visit us at affirmednetworks.com.