

The Battle of SDN vs. NFV

While discussions rage over SDN vs. NFV, the reality is these strategies are complementary – not competitive – and are often deployed together

Executive Summary

For the past 20 years there has been a slow, methodical pace of change in networking where each set of new technologies arrives, surpassing the previous generation in an orderly fashion. But today, two new technologies, Software Defined Networks (SDN) and Network Functions Virtualization (NFV), are poised to truly disrupt that pace by changing networking from physical to virtual. In the enterprise market there appears to be a lot of confusion about the differences between these two, often people compare the two and say one is a better approach than the other. If you are looking for the smack down article that pits these two emerging networking methods against each other in a glorious battle for three-letter acronym supremacy, you've come to the wrong place. Both SDN and NFV have a lot in common with each other. In actuality, the two can coexist in the same network environment and share a lot of the same characteristics and components. What makes them different has more to do with where they started and who is deploying them than anything else. Both methodologies have a similar goal: reduce the cost, complexity and rigid nature of networking, enhancing the physical network with a virtual overlay that is easier to deploy, manage, reprovision and troubleshoot – all with a lower OpEx and CapEx profile.

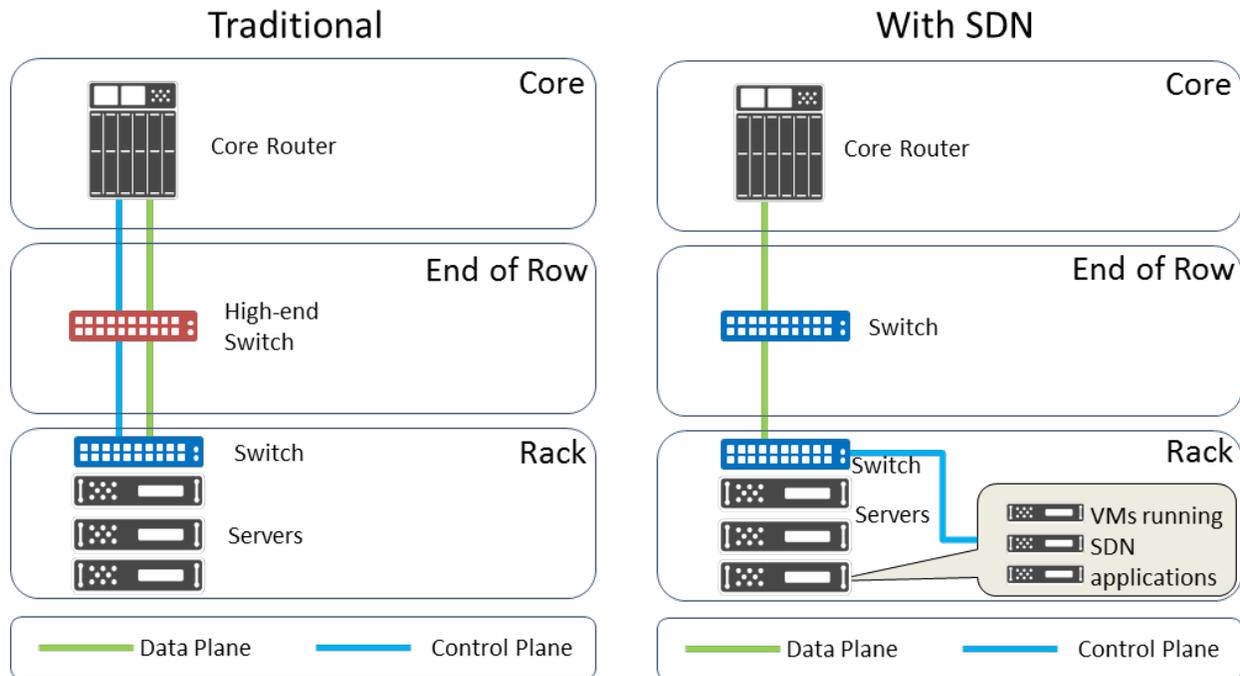
SDN: The Heavyweight

SDN is the undisputed heavyweight in this battle because it receives more press attention and column inches; it's more "top of mind." Beginning all the way back in 2005, it has the first mover advantage for mindshare. Often you see the phrase "born in the campus, matured in the data center" applied to SDN. To some degree the idea that this started in academia might lead people to the conclusion that this is more theoretical than practical, but that couldn't be further from the truth. SDN, while it still has a long hill to climb on broad market acceptance, is quite robust and hyperscale data centers are deploying it today to solve their network challenges.

The goal of SDN is to separate the control plane from the data forwarding plane in the network architecture. This brings more flexibility in how networks are deployed and managed, but most importantly, it allows many of the SDN components to be deployed on industry-standard x86 servers. An SDN Controller, which is a piece of software similar to an operating system, actually manages the aspects of a software-defined network. Companies like [Big Switch Networks](#), [Cisco](#), [HP](#), [IBM](#), [Juniper](#) and [VMware](#) offer SDN controllers. In addition, there are open source controllers also available from projects like [Floodlight](#) and [OpenDaylight](#). ([Click here for our view of OpenDaylight](#)). In SDN the controller runs on a virtual machine (VM) that is hosted on an x86 server, or it can also run on bare metal (x86 or other network device.)

While SDN can reach from the top of the rack all the way back to the service provider, the two most prevalent use cases are connecting remote locations and simplifying networking inside existing data centers. Of those two use cases, the second, simplification, holds the biggest promise for SDN in the sense of mass acceptance and deployment. Today, when customers deploy, move, reprovision or handle just about any task on their network, physically touching a server is not only time-consuming and costly, but physical touches introduce the potential for mistakes. SDN creates a virtual overlay that allows better orchestration of rules-based actions that takes the human element out of the equation and allows changes to happen in seconds, not minutes, hours or days.

Today customers spend a large sum on end of row and top of rack switches, but in the world of SDN, with much of the heavy lifting being done by the SDN controller, lower cost, less feature-laden switches can be used, saving significant money for businesses. There is a wide range of OpenFlow-compatible switches available from the low end (companies like [Pica8](#)) to midrange choices from networking players like [HP](#) or [Dell](#) through high end choices from [Cisco](#) or [Juniper](#). Based both on the scale that they purchase, hyperscale datacenters like Google have actually contracted to have their own switches built based on the particular features that they require/don't require.



NFV: The Challenger

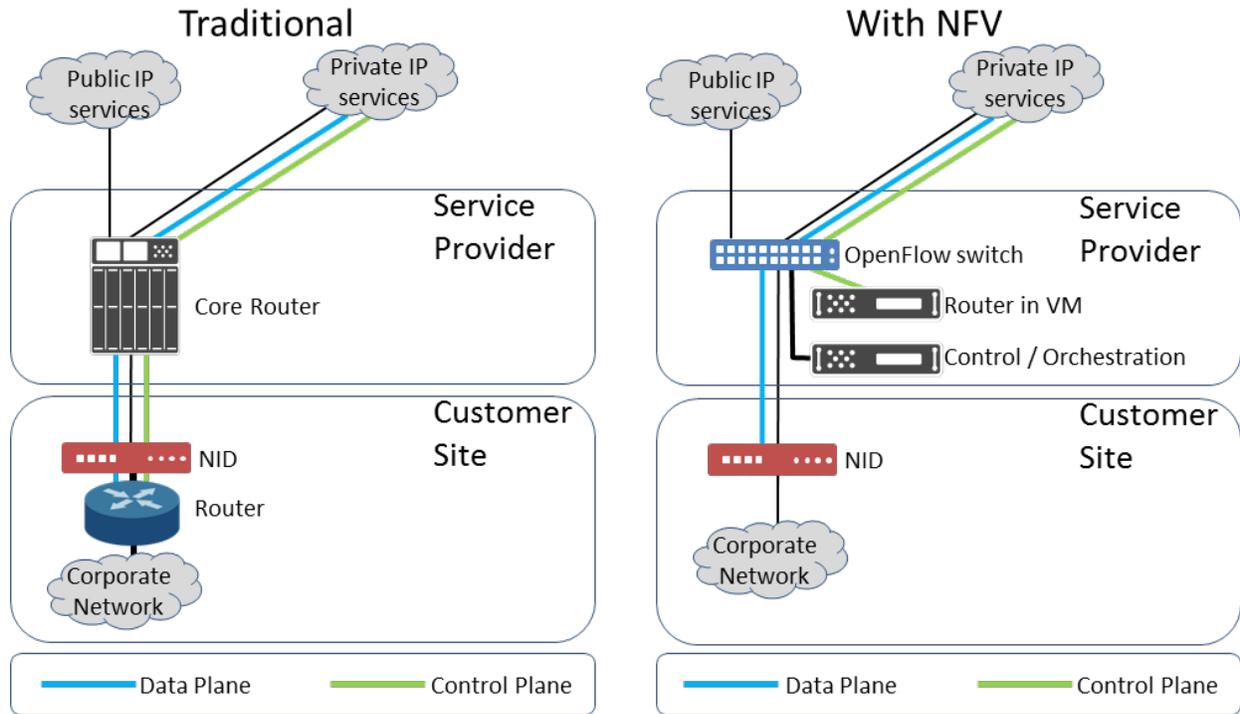
NFV stands for Network Functions Virtualization. Right off the bat, you're probably thinking that description sounds similar to what SDN is doing. Now you can see where the lines blur a bit and the Venn Diagrams start to intersect. Introduced in 2012, this technology direction is rather new, but it has a more streamlined audience and an easier opportunity to really take hold.

NFV is often said to be created by service providers and the focus is clearly where service providers are feeling the most pain. When a service provider makes the connection to a new location, there are several pieces that need to be on-premise for the customer including a managed router and a Carrier Ethernet Demarcation device. (The demarcation device is critical as this creates that "line in the sand" between a customer's network and the carrier's network.)

In addition to these two pieces there may be additional equipment that needs to be installed in order to allow the service provider to monitor, manage and troubleshoot the connection as well as traffic. Normally when a service provider adds another customer or location, there is some set of equipment being installed. Unfortunately most deployments end up with some unusual equipment that needs to be added, so the support matrix, both from procuring and servicing standpoints, becomes rather large for service providers. NFV helps to address this challenge by virtualizing network functions into software applications that can be run on COTS (commercial off the shelf) x86 servers or as virtual machines running on those servers.

With NFV, a carrier may be able to get down to only needing a Network Interface Device for demarcation, allowing the rest of the functions to be housed at the service provider's location and utilizing SDN, of all things, to simplify and manage that network for the customer.

Whether the service provider or the customer is going to manage their network, through NFV, the connection can be moved from a handful of proprietary physical devices down to only one or two physical devices with the rest of the functions being handled virtually via software.



High-level Comparison

	SDN	NFV
Focus	Data Center	Service Providers
Strategy	Split control and data forwarding planes	Replace network devices with software
Protocol	OpenFlow	Not determined yet, does support OpenFlow
Applications run	On industry-standard servers or switches	On industry-standard servers
Customer Benefit	Drives down complexity and cost, increases agility	Drives down complexity and cost, increases agility
Prime Initiative Supporters	Enterprise networking software and hardware vendors	Telecom service providers

Similarities and Differences

As we have said, there is much similarity between the two approaches with both focused on making data communications easier to deploy, manage and change while also driving down the costs. Both approaches focus on replacing proprietary networking equipment with software running either on industry-standard server hardware or on virtual machines hosted on servers. In either case, the fact that the physical platform is being replaced by an x86-driven platform is fundamental to the success. By deploying virtualized networking in VMs on COTS equipment from server OEMs like Dell, HP or IBM, service providers can reduce the equipment variability that they face today, streamlining support and procurement – which has a significant impact on both acquisition costs and operating costs.

Currently the pace of change in networking, combined with the sheer magnitude of platform variety contributes to a financial challenge that is becoming unworkable – both for end customers and service providers. When combined with high costs (because of the proprietary nature of the equipment), longer amortization cycles and higher servicing costs, networking is eating deeper into company budgets (service cost is often equated as a percentage of the total cost of the platform). The ability to grow and scale businesses is being held back by the cost of networking.

In a cloud-focused world where companies on both sides of the demarcation line are being pushed to respond faster, financial barriers can be offset with software-based functionality, allowing for new features to be adapted faster. This gives businesses the velocity to grow, instead of holding them back as they wait for equipment to fully depreciate before making infrastructure improvements.

With both alternatives able to utilize the OpenFlow protocol, there is a natural continuum that allows service providers to offer even more flexible managed IP services, especially for smaller customers who are not in a position to manage their own IT infrastructures.

Where the two approaches differ begins with the organizations behind the initiatives. SDN was originally driven by the need to continually change network infrastructure when testing different solutions. Researchers realized that by separating the control plane from the data forwarding plane and virtualizing all of the connections, they could remove the hard-wired barriers of networks and quickly change structures to suit their needs. SDN was focused mainly from the core down to the server. As new [cloud](#), [big data](#) and integrated applications started to drive more communications between servers, the ability to bring some of the control down to the server level meant that networks could utilize more east-west communication, reducing the north-south traffic and speeding up operations. The key drivers behind SDN are software vendors like [Big Switch Networks](#) and [VMware](#) (through their purchase of [Nicira](#)). Hardware vendors

like Cisco, HP, IBM and Juniper are also involved, but in understanding who has the most to gain from an SDN-driven world, it is clear that hardware players are not always in that camp. ([Here is our take on Cisco's unique challenge for SDN.](#)) While hardware vendors are stepping in to be involved in SDN, they are most likely viewing this as a hedged bet, similar to server OEMs who tenuously approached server virtualization in the early days. Just visiting a vendors' website shows some insight into how deeply the business is aligned with SDN. Search on the term "SDN" or try www.vendorURL.com/SDN and see the outcome. IBM has plenty of content and is clearly embracing the strategy. A search of Cisco.com returns "did you mean ISDN?" in red at the top of the search results.

For NFV, the need was driven by carriers like Deutsche Telekom and BT Group who saw that the cost and complexity was holding back their ability to profitably scale with more customers. Quite often power was becoming a more dominating factor in deployments as well. A change was needed. NFV began in Europe through the [European Telecommunications Standards Institute](#) (ETSI). As a carrier-driven initiative, the primary focus is on how networking functions can be virtualized – mainly along the border between carrier and customer networks – for more efficient and scalable deployments. NFV's principal proponents are carriers like [AT&T](#), [BT Group](#), [Deutsche Telekom](#), [Orange](#), [Telecom Italia](#), [Telefónica](#) and [Verizon](#). While virtualization is a primary driver with this group as well, the key goal is not in the separation of data and control, but more on the standardization of devices. There is more of a primary focus on driving commonality and commoditization at the device level and replacing the physical with the virtual then in changing the routing execution.

Many of the proprietary platforms that providers are deploying today are already built on an x86 architecture (and running on a Linux variant like [Wind River](#) or [MontaVista](#)), so the ability to port that platform to a virtual machine should be relatively straightforward. For equipment manufacturers, however, this removes a revenue component from the equation and also reduces the barrier to switching for the end customer. When a solution is a complete platform, there is a higher switching cost with end customers through each generation of hardware, but when that products becomes a purely software solution, the ability for a customer to move to a competitor's product can happen at any point, not just at the point of platform amortization or platform generational change. In the world of software-focused solutions like SDN or NFV, it is easier for a vendor to "buy out" a competitor's installation as this does not require capital expenditures.

Vendor Impact

What does this mean for the networking world? Clearly there will be a functional shift as vendors need to either augment their portfolios (as many have) with solutions that fit into the needs of these categories.

But there is a strong cautionary tale for vendors to heed. Constituencies on both sides of the line, end customers and service providers, are making their needs crystal clear. Networking today is too expensive, rigid and proprietary; it is holding businesses back across the board. This means that NFV and SDN solutions present a very disruptive force to these incumbent products.

Much of the current incumbent revenue stream is tied to networking hardware which had traditionally been acceptable based on the standard IT workloads. But as technologies like [cloud computing](#) and [big data](#) continue to gain prominence and the world shifts from a desktop client access device to mobile devices like tablets and smartphones, the backend landscape needs to change rapidly in order to keep pace. This will demand more flexible networking solutions that will disrupt the market for the first time in nearly 20 years. Those vendors that can move quickly to address the changing customer needs will benefit and those that try to hang on to “the good old days” may find their share and revenue eroding.

Vendors need to have a strategy that addresses network virtualization (as applicable to their products), but that strategy needs to focus on new products and new deployments, not simply retrofitting older platforms into these new deployment scenarios. SDN needs to be a focus for new networking products targeted for the data center and NFV should be a chief concern for products that are targeted at service providers.

Most of the SDN deployments in the near term are focused on greenfield opportunities and will not be retrofitted into existing data centers. Carriers experimenting with NFV will roll that out with new deployments first; they will not be pulling out existing unamortized equipment, so vendors need to view their products in silos to some degree.

The Future

[Our vision of the future](#) is very different from where we stand today for a few reasons. To begin with, there are fundamental things happening at the platform, software and networking levels that are all in motion today. What the future looks like is open to interpretation, but when you compare where we sit today with virtualization in 50%+ of the server platforms being deployed and huge hyperscale data center that drive web and cloud applications, it is safe to say that the last ten years have brought tremendous change to IT. The next 10 years promise just as much disruption. SDN and NFV will have a place in that reality, but only to the degree that we can find equilibrium with other technology trends that are clearly driving the market as well.

As we said, NFV has a much lower threshold for acceptance based on the dynamics of who is buying it and where it is being deployed. Here are some of the factors that will drive the acceptance or slow the progress of SDN:

- **Hyperscale Data Centers** – These have been some of the first locations to take on SDN, but not every corporate customer is running at that scale. The homogeneity of their applications (and

deployments) makes them exceptionally well-suited for these experiments. It is easy for a [Google](#) or a [Facebook](#) to build their own solution, but as you move down the food chain to regular corporate data centers, the economies of scale won't pay off in the same way. Are these customers a harbinger of the future or an outlier? We believe that it is too early to start translating this trend in hyperscale over onto the rest of the market.

- **Hosted Cloud Solutions** – As more companies rush “to the cloud” or look for hosted cloud solutions outside of their data center, the probability that they undertake a major networking overhaul across their legacy environment goes down. SDN may reign supreme in cloud environments in the near future, but if moving services off premises to the cloud solves many of the current data center challenges of today, is there an appetite for fixing the legacy environments that remain? As a parallel, look at mainframe technology. There is not as much new development in this area, instead it consists mainly of on ongoing operations and maintenance. Do legacy data centers running applications like Exchange, SAP, Oracle and SQL Server meet the same fate?
- **Microservers** – Today's networking problems are exacerbated by the fact that rack density has been on the rise for a long time. 1 and 1/2U servers are now common for mainstream application deployments, meaning that rack density can often be in the range of 40-80 systems per rack – all with multiple physical Ethernet connections per server. But microservers are chassis-based and their fabric allows for as little as 40 to over 750 servers to share a single set of high bandwidth Ethernet connections. If microservers rise in popularity, the internal chassis fabric that is created handles much of the I/O, aggregating Ethernet at the chassis level. This removes many of the existing network issues (i.e. bottlenecks) and reduces the value (and more importantly the ROI) of SDN. If internal east-west traffic is contained inside the chassis and provisioning can be handled virtually, is there a compelling need to deploy SDN?
- **40Gb and 100Gb Ethernet deployment** – There is a phrase in the semiconductor world: “big caches compensate for a lot of inefficiencies.” What about big pipes? In a world where most corporate networks are still dealing with multiple 1GbE connections per server and just beginning to bring 10GbE down to the server, efficiency and bottlenecks are critical issues. But when every server has its own Autobahn and the route back to the core is wide and fast, can't more traffic just go north-south? If network bandwidth goes up by a magnitude of 4x or 10X, does the demand for SDN drop somewhat?
- **Intel disaggregated servers** – [Intel recently announced a strategy](#) that has the potential to change what happens at the rack level. This new strategy relies on silicon photonics to deliver ~6-12GB/s in throughput today (with the potential to take that up to 100GB in the future.) This could remove top of rack and end of row switching from the equation altogether. In doing so, that essentially nullifies much of the need for SDN (however NFV would still be viable in the connection to a service provider). Clearly, as the server landscape is in a state of flux from a form factor perspective, customers may wait to see how things play out on the server side first before diving into the networking as networking sits in the middle (versus servers on the edge) and touches more points.

The Bottom Line

When combining these factors in with the current work being done on SDN and NFV it is clear that there are still many pieces in motion.

For NFV the outlook is more straightforward: it's coming and it is carrier-driven. Just as your cell phone carrier or your cable TV operator will make technology decisions that you will not have a say in, telecom service providers will begin to adopt an NFV strategy and there may be little input that you will be able to provide (but that is not necessarily a negative.) As long they meet the agreed-to SLAs and have an acceptable QoS, there is little ability to influence. IP services, in a simplistic view, are akin to a telephone dial tone. Provided it works as advertised and hits the right metrics, much, if not all, is

forgiven.

For SDN, however, that is a trickier decision. There are still too many emerging options for customers to jump in and take the plunge. On greenfield or opportunistic projects, SDN can be deployed, but in looking at an enterprise-wide standard there are far too many unknowns at this point and we still don't know who the winners and losers will be. Much like standardizing on an OS or a database, the decision to deploy SDN in a wide scale manner carries a significant amount of weight and should not be taken lightly.

What does network virtualization mean for its key constituencies?

- **Networking Hardware Manufacturers** – Need to embrace an SDN strategy that is focused on openness and not simply a new ploy to tie customers into a proprietary stack. This includes supporting multiple controllers and open industry standards like OpenFlow. A key piece of the strategy is understanding the differences so that you are not merely lumping the two together in a “standards” line item on a spec sheet. It is fine to have SDN and NFV support on your spec sheet, but when it comes to products they need to be truly focused on the needs that customers have, not just something to fill a category in the market.
- **Service Providers** – Need to have an NFV strategy in place and begin understanding if/how they can begin deploying NFV in customer locations. With a lower barrier to entry (the support matrix of devices is already long for most) there is an immediate opportunity to leverage the new technology.
- **Server OEMs** – With SDN controllers and applications being deployed on x86 servers, there is a natural opportunity to blend servers with your networking business instead of treating them as functional silos. Your SDN strategy should leverage your strength and footprint in delivering virtualized compute platforms on x86 hardware today (and potentially ARM hardware in the future.
- **SDN Software Developers** – This is an environment of great uncertainty and potential disruption in the market. Despite the interest from both customers and the industry at large, not every SDN company is going to make it in the long run. Most companies will be acquired or exhaust their funding; this means properly scoping funding needs and creating strong industry alliances will be key to the long-term survivability of SDN startups.
- **End Customers** – As we [have said before](#), the SDN market is still emerging and it may be too early to be placing bets. Every enterprise needs to have a virtualized networking strategy; having that strategy can help sort out whether any of the options today meet your needs or whether waiting is still the best path. In terms of NFV, if your service provider is pushing that technology, make sure you understand the impact on SLAs and QoS by moving to the new arrangement.

Important Information About This Paper

Author

[John Fruehe](#), Senior Analyst at [Moor Insights & Strategy](#)

Review

[Patrick Moorhead](#), President & Principal Analyst at [Moor Insights & Strategy](#)

Inquiries

Please contact us [here](#) if you would like to discuss this report and Moor Insights & Strategy will promptly respond.

Citations

This note or paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

Licensing

This document, including any supporting materials, is owned by Moor Insights & Strategy. This publication may not be reproduced, distributed, or shared in any form without Moor Insights & Strategy's prior written permission.

Disclosures

Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness or adequacy of such information and shall have no liability for errors, omissions or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.

©2013 Moor Insights & Strategy.

Company and product names are used for informational purposes only and may be trademarks of their respective owners.