

WHITE PAPER

Designing and Building a Datacenter Network: An Alternative Approach with OpenFlow

Sponsored by: NEC Corporation

Rohit Mehra
January 2012

EXECUTIVE SUMMARY

Server virtualization and cloud computing are changing the face of enterprise computing today. Virtualization enables more efficient use of IT resources and greater levels of IT agility and control. Cloud extends these benefits, allowing IT organizations to reduce their infrastructure complexity, ease staff workload, and more rapidly scale compute resources. Together, these technologies enable organizations to better meet organizational demand and provide greater agility for the enterprise.

Unfortunately, most current network technologies were not developed with the needs of virtualization and cloud computing in mind, and as a result, the network can become a bottleneck to cloud and virtualization deployments. Static topologies require manual intervention to deploy and migrate virtual machines (VMs), which adds cost and burden to IT and hinders the organization's ability to respond quickly to changes in the environment.

OpenFlow is an open source networking architecture that is supported and promoted by the Open Networking Foundation (ONF) and designed to address these shortcomings. An extension of Ethernet, OpenFlow separates the data path and control path, with all networking logic and policies handled by a separate controller. This introduces a new layer of abstraction in networking, analogous to server virtualization, and enables the network to act as a single fabric.

With OpenFlow, the network acts as one "big switch." All logic occurs in software and can be changed as application and network requirements change, allowing for instant and automatic propagation of new policies throughout the network, simplification of network management, and dynamic partitioning of the network to easily handle multitenant environments or traffic segmentation needs as one would using VLANs.

NEC has been a leading supporter of the OpenFlow standard and is the first vendor to bring to market a generally available solution. The NEC ProgrammableFlow family consists of datacenter-grade switches, controllers, and a management console designed to meet the needs of enterprises and cloud service providers. It has already been implemented in several production environments; use cases include hosting/public cloud providers, maximizing value from server virtualization, private/hybrid cloud deployments, network programmability and customization, meeting compliance and audit requirements, and dynamically providing VLAN-type network segmentation.

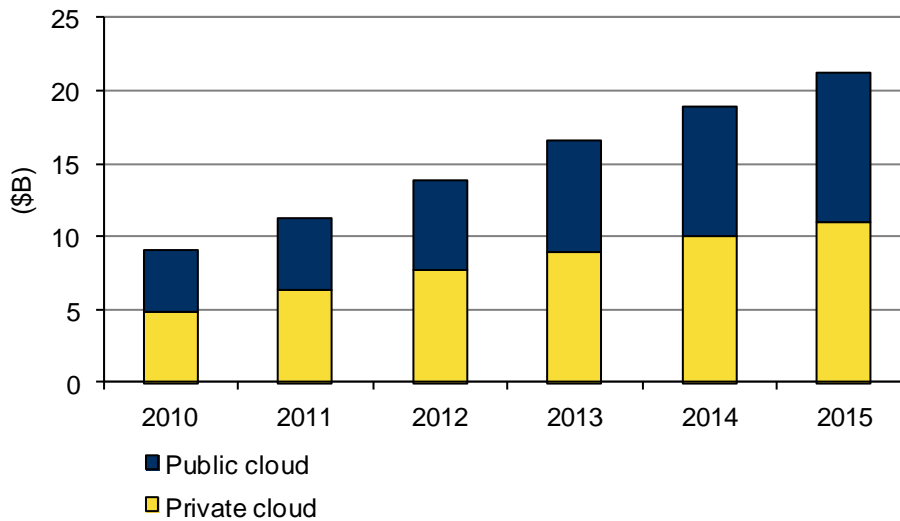
SITUATION OVERVIEW

Cloud Adoption Takes Off

Cloud computing is one of the most important trends today for enterprise IT organizations looking to reduce costs and increase operating efficiency. It allows IT organizations to reduce the complexity of their environment, ease the workload of the internal IT staff, and reduce the type of skills that need to be maintained. IDC sees increasing demand for cloud computing and is forecasting growth across all three cloud models: private, public, and hybrid (see the Appendix for definitions). Looking at hardware investments to support public and private cloud deployments, IDC sees private cloud investments growing to \$11 billion worldwide by 2015, a 17.8% CAGR from 2010, and public cloud investments in third party–manufactured hardware growing to \$8.9 billion by 2015, a 24.0% CAGR from 2010 (see Figure 1).

FIGURE 1

Worldwide Spending on IT Infrastructure Hardware for Public/Private Clouds, 2010–2015



Source: IDC, 2012


Networking Is a Key Investment Area for Cloud Infrastructure

Although server hardware gets a great deal of attention when organizations consider investments they need to make to support a cloud infrastructure, in reality, all aspects of the IT infrastructure must be cloud enabled, including storage, software, and — importantly — networking. Legacy network architectures are designed to support static network configurations and often do not provide the flexibility required to support cloud and datacenter deployments. In fact, IDC sees worldwide investments in datacenter networking equipment to support private cloud deployments increasing to over \$1.2 billion by 2015, a 16.9% CAGR from 2010, and datacenter networking

investments to support public cloud growing to over \$1.3 billion by 2015, a 23.3% CAGR from 2010 (see Figure 2).

FIGURE 2

Worldwide Spending on Datacenter Networking for Public/Private Clouds, 2010–2015



Source: IDC, 2012

Tiered Networking and Heterogeneous Environments Create a Network Hodgepodge

To ensure their most efficient operation, networks, storage, and servers need to work together seamlessly and efficiently. Unfortunately, network architectures in many enterprise datacenters consist of a hodgepodge of architectures, protocols, and devices.

The network hodgepodge is due to a variety of factors. Many companies pursue corporate growth through mergers and acquisitions, and IT is tasked with integrating datacenters with disparate network topologies and architectures as the newly acquired companies are brought into the corporate fold. In other cases, IT infrastructures have grown through waves of capital purchasing, and as new protocols emerge, organizations modify their strategy, or vendors change product lines, enterprises do not always purchase the same equipment the next time around.

Server Virtualization Moves Mainstream

Virtualization is another key trend shaping the enterprise datacenter today. It reduces physical server sprawl and thus reduces up-front capital costs and operating costs in management and power and cooling and extends the life of the datacenter. Furthermore, by introducing a layer of abstraction between the OS/application layer and the underlying hardware, virtualization provides a greater degree of flexibility, agility, and control over the infrastructure. Many datacenters have virtualized a significant portion of their workloads. IDC research shows that as of the end of 2011, the installed base of virtual servers was roughly equal to the installed base of physical servers and will continue to grow dramatically in the years to come.

Further, many IT organizations have neither the budget nor the stomach for replacing their entire network infrastructure outright, so they choose instead to deploy the new equipment in a segment of their network — for example, in the network core or in specific datacenters.

Whatever the driver, many enterprises currently have "pockets" of different architectures and topologies and are left with highly heterogeneous networking environments. This introduces challenges to organizations as they look to operate a seamless network across their entire IT infrastructure and complicates the process of introducing new equipment into the environment as it must interoperate with legacy networking infrastructure. Equipment from different vendors uses different management tools, operating protocols, and alerts/error messages, and network managers need to develop proficiency in each. Further, each vendor's equipment typically requires the use of that vendor's network management system because of the proprietary nature of many vendors' management information schemes. The process of dealing with day-to-day operating issues and resolving network issues becomes more complex and time consuming because of the different systems involved, and the problem is compounded when adds/changes/moves need to be made.

The process of dealing with day-to-day operating issues and resolving network issues becomes more complex and time consuming because of the different systems involved, and the problem is compounded when adds/changes/moves need to be made.

Current State of Enterprise Network Architecture Inhibits Agility, Flexibility

Unfortunately, the current state of many enterprise networks inhibits the ability to support the demands of virtualization and cloud computing and, at a deeper level, IT's ability to rapidly respond to line-of-business demands. The result is that the network acts as a bottleneck to enterprise agility and is a limiting factor to the organization's ability to scale and adapt to changing market conditions.

The current state of many enterprise networks inhibits the ability to support the demands of virtualization and cloud computing and, at a deeper level, IT's ability to rapidly respond to line-of-business demands.

Most corporate networks are not designed to be aware of virtualization or cloud computing. They assume a static application topology and treat all servers and clusters the same, regardless of whether they are physical or virtual. Manual intervention is required to move virtual machines, and most IT departments rarely migrate virtual machine workloads. Most current datacenter architectures were designed to enable heavy loads in "north-south" traffic (between clients and servers) and are less well adapted to server-to-server traffic that is more emblematic of virtualization and cloud computing. Tiered networks and heterogeneous environments exacerbate the problem and further hinder network managers' ability to introduce flexibility into the network.

The Need for an Alternative Approach

As IT organizations struggle to enable greater organizational agility with support for virtualization and cloud computing, many are realizing that their network is at the breaking point. Traditional fixed network approaches are simply not designed to meet the current needs of a virtualized environment and are not flexible enough to support the business' changing demands. Instead, IT organizations need a new approach in which the network can be treated more dynamically and in which VMs and other resources can be quickly and flexibly introduced, moved, or modified as needs change, without requiring manual intervention to reconfigure the network. They need a network that is flatter with fewer layers, supports easy migration of virtual machines,

Traditional fixed network approaches are simply not designed to meet the current needs of a virtualized environment and are not flexible enough to support the business' changing demands.

supports multitenant environments, and has low intra-datacenter latency. With this capability, the network could become a corporate asset that enables organizational agility rather than a liability.

THE OPENFLOW SOLUTION

OpenFlow is an open source protocol pioneered by Stanford University faculty with funding from the Clean Slate Lab hosted at Stanford. The Clean Slate Lab was cofounded by Stanford University, Deutsche Telekom, and NEC Corporation.

OpenFlow is now managed and overseen by the Open Networking Foundation, a nonprofit organization hosted at Stanford University and backed by leading networking, software, telecommunications, and hardware companies around the world, including NEC. OpenFlow is owned by the community with the goal of improving research and innovation in networking and allowing networking equipment vendors to integrate it into their commercial offerings.

The OpenFlow Switching specification was created in 2008 and is based on Ethernet technology. With OpenFlow, network logic moves into the application stack, decoupling network software from the underlying hardware. Packet transferring and routing control functions are done outside the switches themselves, which enables advancements in routing control, network virtualization, and visualization.

With OpenFlow, network logic moves into the application stack, decoupling network software from the underlying hardware.

In a standard network, both the packets being routed (the data path) and the routing decisions (the control path) are handled by the router or switch. In the OpenFlow approach, these functions are separated. The router or switch handles the data path, and a separate, programmable controller handles the control path. The switch and controller communicate via the OpenFlow standard. This additional layer of abstraction introduces greater flexibility into the network and simplifies management, provisioning, and configuration of network devices. It enables greater virtual machine mobility, the potential for enhanced security capabilities, and support of next-generation IP-based mobile networks.

OpenFlow also enjoys the benefits of an open architecture approach. While the Ethernet standard is well established, each vendor's implementation is somewhat different, and commercial switches and routers typically do not provide an open software platform supporting innovation or customization. An open systems overlay over Ethernet, OpenFlow lowers the barrier for entry for new ideas, allows customers to avoid vendor lock-in, and could help increase the rate of innovation in network infrastructure.

An open systems overlay over Ethernet, OpenFlow lowers the barrier for entry for new ideas, allows customers to avoid vendor lock-in, and could help increase the rate of innovation in network infrastructure.

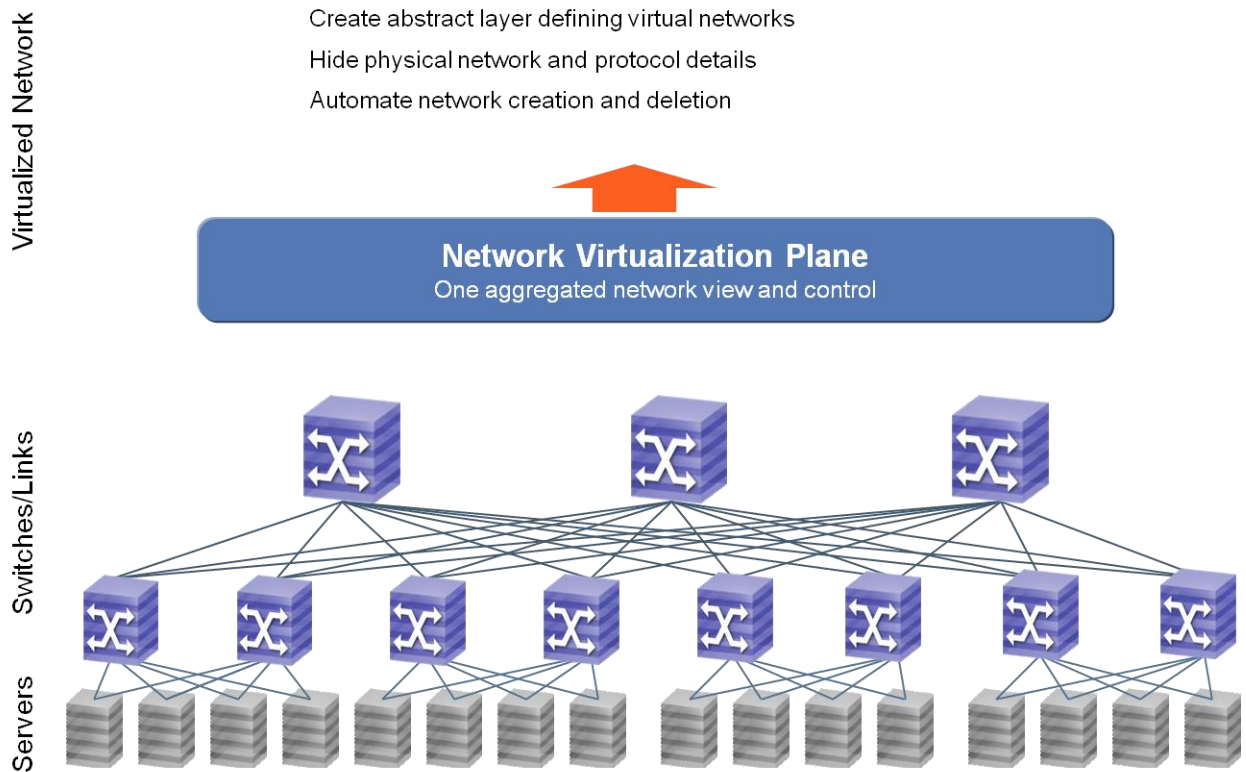
NEC OpenFlow Fabric

The NEC ProgrammableFlow Network Architecture and Product Family essentially provide an enterprise-grade datacenter networking fabric. Designed as a simplified architecture for datacenter, private cloud networks, and the public cloud, ProgrammableFlow leverages the OpenFlow protocol to create Software-Defined Network (SDN) virtualization, allowing customers to more easily deploy and manage multitenant network infrastructure. It provides network virtualization for datacenter

managers by abstracting the network definition and control layer, enabling users to create "virtual networks" and automating the creation and deletion of networks while hiding the physical network and protocol details. This abstraction layer is demonstrated in Figure 3.

FIGURE 3

NEC ProgrammableFlow Enables Network Virtualization



Source: NEC, 2012

The ProgrammableFlow family consists of the following primary components:

- ☒ **ProgrammableFlow Controller.** ProgrammableFlow Controller software brings virtualization to enterprise networking. Acting as a control center, the ProgrammableFlow Controller deploys switching logic to each of the ProgrammableFlow- and OpenFlow-enabled switches in the network. This allows datacenters to deploy, control, monitor, and manage multitenant infrastructures from a single console. Physical server and virtual machine network interfaces can be controlled via this console, enabling rapid scale-out of new applications without requiring network managers to manually reconfigure the network. The ProgrammableFlow Management Console provides end-to-end visualization of the physical network and each virtual tenant network.

- ☒ **ProgrammableFlow PF5240 Switch.** The ProgrammableFlow PF5240 Switch from NEC is an Ethernet-based switch that can integrate into legacy Ethernet environments while simultaneously functioning as an OpenFlow Switch. It was the first OpenFlow Switch to be generally available to customers worldwide and has 48 GbE ports and 4 10GbE ports. It supports line rate multilayer switching, maintaining up to 160,000 network "flow entries" or units of OpenFlow communication.
- ☒ **ProgrammableFlow PF5820 Switch.** The latest member of the ProgrammableFlow family, the NEC PF5820 Switch supports 10/40GbE and is designed to support high-performance, low-latency, energy-efficient OpenFlow networks. It provides 48 10GbE SFP+ ports plus 4 QSFP+ ports that operate at 40GbE or as 16 additional 10GbE ports, providing a total of 64 10GbE ports in one switch. It supports up to 96,000 flow entries.

ProgrammableFlow is designed to allow datacenters to streamline network management through greater levels of automation while lowering operating costs and reducing the time required to deliver network services. Its aim is to allow customers to more easily upgrade and add functionality to their networks while protecting their network investment.

Advantages of Software-Defined Networking

SDN, a key tenet of the ProgrammableFlow approach, promises to bring new levels of flexibility, management, and control to networking. SDN is analogous to the changes brought about by server virtualization in which IT managers can decouple applications and their environments from the underlying hardware, introducing greater levels of freedom and flexibility in the server IT infrastructure. Similarly, by decoupling network logic and policies from the underlying switching hardware, SDN brings new flexibility into the networking environment. Logic and policies can be defined, changed, and modified on a regular basis, all from a central location, and propagated throughout the network. Furthermore, the networking logic at the controller level is truly programmable, meaning that new services can be defined, innovations can be introduced, and the network can be customized to suit the needs of the organization as needed. This programmability introduces even greater levels of flexibility, innovation, and control than those achievable by simply centralizing policies into a central management dashboard.

The networking logic at the controller level is truly programmable, meaning that new services can be defined, innovations can be introduced, and the network can be customized to suit the needs of the organization as needed.

Key Benefits of NEC's OpenFlow Approach

The OpenFlow approach to networking that is incorporated into NEC ProgrammableFlow is intended to provide a number of benefits to an organization, including:

- ☒ **Simple.** Multitenant virtual networks can be deployed as easily as VMs in a server infrastructure setting. Network and application policies can be implemented for centralized network management and control, and OpenFlow eliminates the need for spanning trees or other distributed protocols, reducing network complexity and enabling organizations to unlock trapped network capacity.
- ☒ **Open.** With the OpenFlow technology, customers can create multivendor configurations of switches, controllers, and NICs. OpenFlow can be deployed in heterogeneous network topologies, increasing network resiliency and capacity.

- ☒ **Scalable.** OpenFlow scales from single switches to full-scale datacenter fabrics. Policy-based appliance integration enables implementation on high-performance network devices.
- ☒ **Fast.** ProgrammableFlow currently supports 10GbE switches. Hardware forwarding, quick convergence times, and network load balancing all contribute to the speed of the solution.

These benefits are designed to enable companies to reduce their overall operating and maintenance expenses as well as capex costs for new network equipment. Organizations can increase server and network utilization, while greater agility in deploying networking services and application resources enables organizations to move faster to meet changing market needs.

OpenFlow Use Cases

IDC sees a number of scenarios for OpenFlow capabilities in commercial deployments, including:

- ☒ **Hosting/public cloud providers.** With its support of multitenant network environments, all managed from a centralized interface, OpenFlow is a strong fit for hosting or public cloud providers. These providers can deploy and implement network resources for individual customers and rapidly reprovision or scale them up or down as required in a highly automated manner, without requiring manual intervention.
- ☒ **Maximizing value from server virtualization.** Network limitations often act as bottlenecks to organizations getting the most from their server virtualization efforts. By acting as a virtual network fabric, an OpenFlow network provides seamless integration into a virtual server environment, enabling servers and VMs to be provisioned, migrated, and decommissioned without requiring network reconfiguration. Network and security policies follow virtual machine migrations automatically.
- ☒ **Private/hybrid cloud deployments.** Just like hosting or public cloud providers, enterprises deploying private or hybrid cloud implementations must also manage multitenant environments and need to scale network resources up or down in each of those environments quickly and easily with minimal impact on IT staff time. The OpenFlow approach is ideally suited to satisfying these requirements.
- ☒ **Network programmability and customization.** Decoupling the data path and control path lets organizations more easily introduce changes into their network and customize it to suit their particular purposes. Customers can replace custom-built networking appliances with general switching equipment in which the "heavy lifting" logic is performed by the controller. For example, many organizations install special hardware in front of their network infrastructure to protect against denial of service attacks. This equipment is unutilized 99%+ of the time. Instead, organizations can include denial of service logic in their controller and use their general switching equipment to route suspect traffic to a cleaning room, eliminating the need to spend additional capex dollars on these devices.

- ☒ **Security and compliance requirements.** Managing security and access control across large enterprise networks is cumbersome and time consuming, and it is often difficult to implement consistent and current policies across the entire network to ensure compliance with regulatory and audit requirements. With OpenFlow, it is simple and straightforward for network administrators to define networkwide policies in the central controller and propagate them throughout the entire network with the push of a button.
- ☒ **Providing dynamically segmented networks (such as VLANs).** With OpenFlow, administrators can easily partition the network and provide specific users with access to their own isolated network, just as VLANs do. But unlike most VLAN technology in which networks are static, OpenFlow enables administrators to set up and tear down these segmentations dynamically in software, providing greater agility for the organization.
- ☒ **Resiliency and performance in multipath networking.** Network demands are becoming more complex in the era of cloud computing and server virtualization. Many organizations are implementing multipath networks to address these requirements, but traditional network approaches are not well equipped to handle multipath networks. OpenFlow supports multipath networks with custom-defined policies and continuous updates based on network resources and traffic conditions. It maximizes the use of network resources and multiplies the bandwidth available within the network.

CASE STUDY: GENESIS HOSTING SOLUTIONS

Genesis Hosting Solutions is a virtual infrastructure hosting company. It provides highly flexible, customizable hosting services targeted toward service providers and enterprises. Its "build your own cloud" approach enables customers to build and provision customized, highly available virtual machine clusters by assembling resources consisting of CPU, memory, disk, and networking components in a multitenant environment.

Genesis needed to scale networking capacity to support its growing business and was looking to add additional datacenter capacity. In particular, it was facing two key issues:

- ☒ The staff burden and time required to implement complex network reconfigurations was becoming expensive.
- ☒ The spanning tree protocol support across its networking equipment was becoming a roadblock.

Genesis needed a solution that not only would allow it to add new services without requiring hardware upgrades and improve its network service-level agreements but also would not introduce additional network complexity. The solution needed to handle high performance and support a multitenant infrastructure environment.

After evaluating several alternatives, Genesis selected NEC ProgrammableFlow. According to Genesis, one of the chief attractions of NEC ProgrammableFlow is that

the entire network acts as one "big switch" in which any changes or policies that are implemented are automatically propagated throughout, simplifying configuration and policy management and eliminating the need to maintain complex spanning trees.

Another attraction was ProgrammableFlow's support for legacy Ethernet environments. Genesis implemented ProgrammableFlow at its network core but still maintained legacy equipment at the network edge. With its open architecture and support of Ethernet standards, ProgrammableFlow is "ideal" for Genesis' multivendor environment.

Genesis sees a number of benefits to its adoption of ProgrammableFlow, including improved efficiency in its network; unlike traditional Ethernet, OpenFlow allows Genesis to use every port in the network for application traffic (and not just connectivity), which helps it use network resources more efficiently. Genesis is also able to save network management staff time; time that would have previously been required to make network and policy changes can now be spent on more strategic activities. Overall, Genesis estimates that it has reduced network administration by 100 hours weekly, reduced IP address usage up to 60%, and supported its 99.999% service offering with the aid of ProgrammableFlow self-repair.

Genesis plans to continue scaling and building out its ProgrammableFlow infrastructure. Its initial implementation included 1GbE ProgrammableFlow switches, and the company is planning to implement new 10GbE switches in the future.

OPPORTUNITIES/CHALLENGES

IDC sees a number of opportunities and challenges for OpenFlow as a standard and for NEC as it brings its commercial OpenFlow offering, ProgrammableFlow, to market.

Opportunities include:

- ☒ **For customers: more easily embracing virtualization and cloud computing and integrating networking into IT infrastructure.** Servers and storage are important, but their performance is only as good as that supported by the underlying network architecture. Customers that rearchitect their networks to take full advantage of OpenFlow can achieve the many benefits associated with virtualization and cloud computing, including increased flexibility in their networks, simplified management, and faster deployment cycles.
- ☒ **For enterprise IT: demonstrating innovative value-add to the organization.** This is an opportunity for IT to "look good" and to demonstrate its value to the business by providing improved efficiencies that will better support business agility and upcoming IT initiatives.
- ☒ **For NEC: establishing leadership in a rapidly emerging networking area.** The network equipment market is highly competitive, with vendors competing on the need to bring to market innovative new technologies and to develop solutions that reduce companies' total cost of ownership and drive attractive return on investment. By coming to market quickly with an important new networking approach, NEC has the opportunity to establish itself as a leading-edge vendor in this space and gain market share.

Challenges include:

- ☒ **Early-stage market.** OpenFlow is at an early stage in the market and is still evolving. At present, many vendors have yet to jump on the bandwagon, deployments are still relatively few, and things may continue to change before the dust settles. Many customers have yet to be educated about OpenFlow, and not all members of the ecosystem (e.g., VARs and service providers) are fully on board. This may make it more challenging for customers that want to implement OpenFlow to find the appropriate resources and assistance to do so.
- ☒ **Need for vendor adoption.** Many of the leading vendors in networking have yet to adopt OpenFlow in their commercial offerings. It is not clear to what extent they are embracing it internally and putting R&D dollars behind the approach, yet many customers look to these vendors for solutions, both for the validation that broader market adoption provides and to ensure their adopted technologies will be supported for years to come. Adoption by other leading vendors besides NEC will be critical to the success of OpenFlow and its market adoption.
- ☒ **Challenges with the open approach.** A related point stems from the fact that OpenFlow is an open project. Open source projects create a unique set of challenges for vendors as they seek a business model that enables them to appropriately differentiate their portfolios and monetize their offering while still holding to the open model. The degree to which vendors embrace this model will likely have a profound impact on the adoption of OpenFlow.

CONCLUSION

Organizations are moving quickly to implement virtualization and cloud computing. While these technologies provide greater agility, management, and efficiency in enterprise IT infrastructures, they also introduce a number of new challenges to the network that traditional approaches are not designed to address.

OpenFlow is a recent open source technology that has emerged to address these challenges. An extension of Ethernet in which the control path is separated from the data path and is handled by a separate controller, OpenFlow enables the networking logic to be handled in software and introduces a number of efficiencies and new capabilities to handle complex, multitenant network environments.

NEC is a strong contributor to the OpenFlow standard, and its ProgrammableFlow is one of the leading commercial-grade implementations of the technology. ProgrammableFlow has already been implemented in production environments and supports a number of enterprise use cases, including hosting/public cloud providers, virtualization, private/hybrid cloud deployments, network programmability and customization, and compliance and audit requirements.

APPENDIX: THREE CLOUD MODELS

Cloud computing encompasses three primary models:

- ☒ **Public.** In public clouds, the deployment is open to a largely unrestricted universe of potential users, provided by a service provider, and targeted to be used by a market and not a single enterprise.
- ☒ **Private.** Private clouds are designed for access restricted to a single enterprise and are used as an on-premise shared resource and not a commercial offering.
- ☒ **Hybrid.** In hybrid cloud deployments, a portion of the infrastructure is deployed in the service provider's cloud and a portion is deployed at the customer's premises. This way, elements that require high degrees of scalability can be deployed at the service provider while still providing an element of security and control by locating certain data and functions at the user's site.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2012 IDC. Reproduction without written permission is completely forbidden.