IDC TECHNOLOGY SPOTLIGHT

Optimizing the Datacenter Network for Improved Scalability, Orchestration, and Automation

April 2014

Adapted from *Cisco's Approach to Professional Services for Its Application Centric Infrastructure (ACI)*
by Leslie Rosenberg and Brad Casemore

Sponsored by Cisco Systems

The four pillars of computing — cloud, mobility, social, and analytics — are driving new levels of network innovation in datacenter networks. These forces are now buffeting the datacenter along with virtualization and the Internet of Things (IoT), resulting in sweeping changes in traffic patterns that expose the limitations of traditional networks and their operational models. To become a resource rather than a bottleneck to overall datacenter performance, the network must deliver not just exceptional performance and scalability but also unprecedented automation and orchestration that can yield agility, flexibility, and service velocity. This Technology Spotlight examines these key trends and discusses the role that Cisco's Application Centric Infrastructure (ACI) plays in addressing these ongoing challenges for enterprise IT and network managers.

Introduction

The current evolution of computing is prompting enterprise datacenters and their associated networks to adopt new approaches. The four pillars of computing — cloud, mobility, social, and analytics — are now the key trends driving new levels of network innovation that will determine the design and operation of datacenter networks over the next several years.

Like other types of datacenter infrastructure, the network exists to support and deliver enterprise application workloads that are critical to business success. During the client/server era of computing, the network fulfilled that role admirably both architecturally and operationally. In recent years, however, new types of workloads and traffic patterns have made new demands on datacenter infrastructure in terms of responsiveness, user satisfaction, and the ability to roll out cutting-edge applications that depend on enhanced performance.

While other aspects of datacenter infrastructure have changed to meet those challenges, progress in network infrastructure has been slower for a variety of reasons. However, momentum is shifting, and now network managers — with the backing of IT staff — need to step up to the challenges posed by these fundamental changes in the enterprise compute model by adopting new models for network transformation.

Key Trends Reshaping Datacenter Infrastructure

The forces now buffeting the datacenter and network include virtualization, cloud, mobility, big data, social computing, and the Internet of Things (IoT). Many enterprises have virtualized major portions of their application workloads to improve the utilization rates of compute and storage resources and reduce opex.

Virtualization delivered the business benefits of server consolidation and paved the way for cloud computing (see Figure 1). In addition, the proliferation of mobile devices and accompanying application traffic brought sweeping changes in traffic patterns that extended from the mobile endpoint to the datacenter.

Figure 1

Worldwide Server Virtualization Forecast

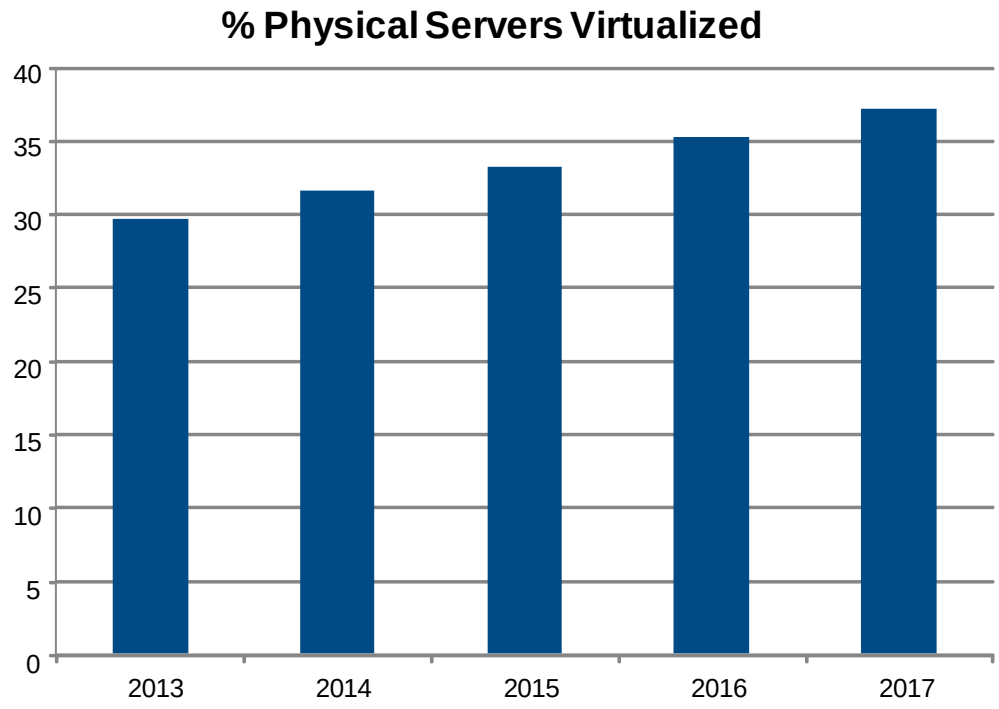

Source: IDC

Virtualization exposed many of the limitations of traditional networks and their operational models. In a virtualized environment, where applications can be spun up as virtual machines (VMs) in a matter of minutes and moved within seconds, IT departments increasingly recognized that the network — built for large-scale, but static IP networks with largely stationary endpoints — had become an impediment to operational agility and positive business outcomes.

The disconnect was clear: VMs could be provisioned almost instantaneously, whereas manual, CLI-based configuration of network devices could take days or weeks. In addition, automation tools sped compute and virtualization processes, while the network was left behind. Virtualization, however, was only a part of the story. Big data and social networking are also contributing to increased traffic and changing traffic patterns. The Internet of Things is destined to bring further change. Finally, cloud — with its private, hybrid, and public variants — is recasting how and where workloads are deployed and how fast data is utilized.

The Benefits of New Approaches to Datacenter Networking

To become a resource and not a bottleneck to overall datacenter performance, the network must be based on an agile, flexible model rather than a rigid and comparatively outdated model. Datacenter IT teams require an advanced network designed around the needs of applications and application developers and one that is capable of separating application policy from operational requirements. To accomplish these goals, the network must:

- Be automated and provisioned with greater speed
- Be managed on a programmatic basis
- Offer real-time telemetry and visibility
- Deliver consistent performance at high scale
- Integrate seamlessly with industry-standard cloud orchestration platforms

With its multilayered architectures, time- and resource-consuming CLI configurations, and siloed operational model, the traditional network proved incapable of adapting to these needs. Even first-generation software-based LAN emulation solutions attempted only to virtualize the existing traditional design without shifting the focus to how to best deploy applications. In addition, they added further complications by creating separate overlay and underlay management domains, offering limited scale at high performance and limited interoperability with existing physical infrastructure.

Companies ranging from cloud service providers to enterprises of all sizes seek a datacenter network that delivers not just exceptional performance and scalability — though those remain important considerations — but also unprecedented automation and orchestration that can yield agility, flexibility, and service velocity. This approach will be especially important for emerging cloud environments (see Figure 2).

Figure 2

Enterprise Priorities for Private Cloud

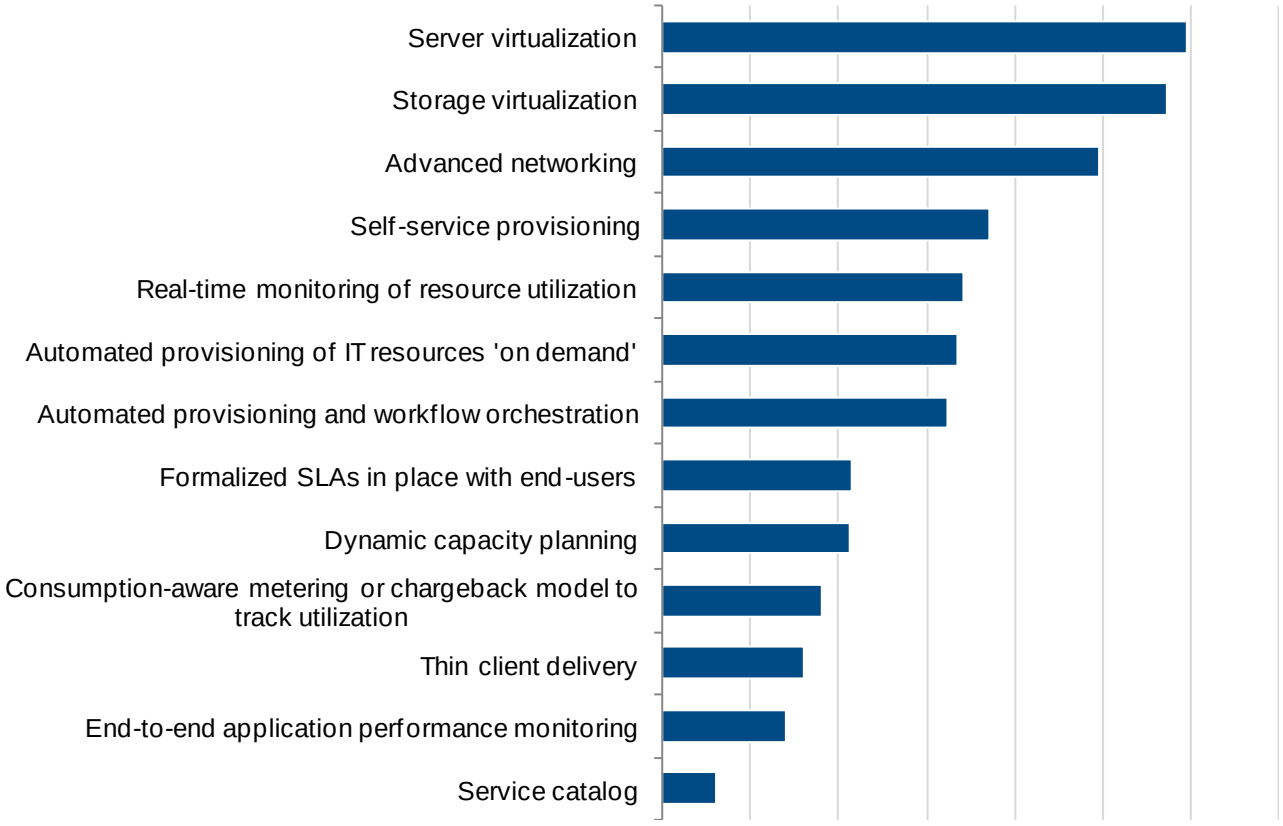

n = 801

Source: IDC's *Cloud Computing Survey*, Spring 2012

In this context, the network needs to be closely aligned with the requirements of application workloads, which run increasingly in and between servers (east-west) rather than from client to server (north-south). Network infrastructure needs to be robust at its architectural base while providing support for industry-standard hypervisors, protocols, and platforms such as OpenDaylight controller and OpenStack orchestration. It also needs to support open APIs. Software-defined networking (SDN), with its decoupling of the control plane and data plane, arose as a response to this set of problems. However, at present, SDN remains a work in progress across the industry.

Considering Cisco's Application Centric Infrastructure

Cisco's Application Centric Infrastructure (ACI), which includes the Application Policy Infrastructure Controller (APIC), the Nexus 9000 family of datacenter switches, and a set of southbound and northbound APIs, represents an alternative to SDN as currently conceived.

Cisco unveiled ACI in November 2013 and contends that it offers significant advantages over existing software-only approaches by providing a hardware and software solution that brings together physical and virtual infrastructure. The result is a platform intended to provide scale, performance, flexibility, and automation for datacenters encompassing both physical and virtual workloads while also enabling real-time visibility and security for network operations.

Key components of Cisco ACI include:

- Application Policy Infrastructure Controller — a policy controller that assigns application-related service levels and access privileges
- Nexus 9000 family of modular and fixed 10/40GbE switches
- An optimized version of the NX-OS that runs in two modes — a standalone mode and an enhanced mode for ACI deployments
- A rich ecosystem of partners across hypervisors, orchestration, DevOps and automation, systems management, L4–7 services, and network monitoring

Cisco's ACI has been constructed around a policy model designed to capture application requirements and automate deployment of those applications across the network. This approach is what Cisco calls a declarative management model, which the company believes to be eminently suited to aligning workloads and datacenter infrastructure while serving the burgeoning requirements of new workloads and cloud computing.

Declarative management is the voluntary cooperation of individuals or agents that publish their intentions via commitments to each other. The intentions are abstract; thus, for example, an application policy would state its requirements, and the underlying infrastructure (e.g., datacenter switches) would interpret how best to satisfy those requirements based on its own capabilities.

According to Cisco, by distributing complexity to the edges of the infrastructure, the approach has excellent scale characteristics. Additionally, its use of abstract policy allows broad interoperability across devices without limiting a vendor's ability to expose a differentiated feature set. Abstract policies also give application administrators a self-documenting, portable way of capturing their infrastructure requirements.

As part of ACI, Opflex is a new, open southbound protocol facilitating declarative management of network infrastructure. It is driven by Cisco with industry support from a number of vendors. Cisco posits that today's networks and first-generation SDN implementations are "high-touch, micromanaged environments," which result in problems involving scaling, coping with failures, and interoperability.

Traditional SDN, as defined by the Open Networking Foundation (ONF), has a policy manager and a decoupled control plane (SDN controller), which uses protocols such as OpenFlow and OVSDB to program flows on physical and virtual switches. Alternatively, Cisco's ACI features the APIC controller as the policy manager, but it relies on intelligent devices implementing a control plane and a data plane to enforce policies.

Opflex is described by the company as a flexible, extensible policy resolution protocol for declarative management of any datacenter infrastructure. Cisco plans to standardize Opflex and is now steering it through IETF approval. The company has invited all vendors to participate in its development as an open protocol and plans to offer an open source agent that any third-party vendor can adopt. Cisco plans to offer support for Opflex through its Nexus 7000, ASR9000, ASA, and Sourcefire products, allowing them to integrate with APIC. Opflex will also be supported by a number of partners, including Microsoft, Red Hat, F5, Canonical, and Citrix.

Opflex is positioned as existing within an open API ecosystem that spans southbound and northbound protocols. These APIs allow integration with a rich ecosystem of partners, including those that offer automation tools (Opscode Chef, Puppet, CFEngine, Python), hypervisor management (Oracle, Xen, Red Hat, VMware, Microsoft), enterprise monitoring (IBM Tivoli, Splunk, CA, NetQoS, NetScout), systems management (IBM Tivoli, CA, BMC, HP), network services (Citrix, F5, Palo Alto Networks, Embrane), and orchestration frameworks (OpenStack, CloudStack, VMware, Cloupia).

The OpenDaylight open source SDN controller platform also will provide support for Opflex, accommodating it as a southbound interface from OpenDaylight to network switches. Cisco is working in concert with other vendors to deliver an OpenDaylight plug-in architecture that supports Opflex.

Challenges

Cisco faces a number of challenges associated with ACI and the APIC controller and with its southbound and northbound protocol strategy. Challenges include the following:

- Uncertainty over how various market segments such as Web-scale, traditional, and cloud service providers will solve their network infrastructure challenges
- Competition from new and traditional networking competitors and alternative architectures in the form of SDN, overlay-based network virtualization, and network disaggregation
- The perception that Cisco doesn't have a fully open architecture (To the extent that customers believe that Cisco is not offering an open, interoperable approach to addressing their problems — thereby protecting their existing investments and giving them flexibility to adapt to evolving technologies and changing needs — they will be wary of committing to ACI.)
- Building a strong ecosystem of committed technology and channel partners (Ecosystems are always important, but they have become critical to success in the current era of virtualization, cloud, mobility, social media, Big Data, and IoT. Cisco must ensure that ACI receives strong support from a wide array of capable and well-regarded technology partners. It will be essential for Cisco to clearly and persuasively communicate how ACI integrates with cloud orchestration systems and industry-leading automation tools.)

Conclusion

Enterprises and service providers are reassessing how their datacenter infrastructures, including their datacenter networks, must be recast to support new and existing application workloads, as well as growing and changing traffic patterns. Virtualization, cloud, enterprise mobility, Big Data, social media, and the Internet of Things are emerging as major transformative factors that must be accommodated. At the same time, these companies must build an infrastructure that is equally adept at handling physical workloads — either traditional client/server applications or low-latency applications — as well as a growing number of virtualized workloads.

As such, the need has arisen for a network infrastructure that provides not only the traditional values of performance, reliability, and scale but also new capabilities such as support for automated provisioning, programmatic management, and integration with industry-standard cloud orchestration platforms.

Various approaches have been proposed to deliver this next-generation datacenter infrastructure. SDN is a notable candidate, offering a decoupling of the data and control planes, but network disaggregation and overlays for network virtualization also have been advanced as solutions. Above all, companies have begun to recognize that the policy layer, which provides the interplay between business-critical application workloads and how they are accommodated by the underlying network infrastructure, is a key consideration. Industry-standard southbound protocols such as Opflex, will be integral to how policy layers effectively communicate and enforce top-down application requirements on the underlying infrastructure. To the extent that Cisco can address the challenges described in this paper, IDC believes the company is well positioned for success in the datacenter automation market.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com