Analyze the Future

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

WHITE PAPER

Ethernet Fabric for the Cloud: Setting the Stage for the Next-Generation Datacenter

Sponsored by: Brocade Communications Systems Inc.

Lucinda Borovick
March 2011

EXECUTIVE SUMMARY

A new era of IT innovation is under way in the datacenter, and it is aimed at two separate but related goals. Specifically, enterprises are looking to accomplish the following:

- Quickly and efficiently drive revenue growth in today's fast-paced, information-centric environment
- Reduce the cost and complexity of today's current IT infrastructure

Server, storage, and network virtualization will be at the core of most datacenter buildouts during this era. Virtualization combined with IT's hopes for cloud computing will significantly influence the way organizations deploy and manage datacenter networks in today's fast-paced, information-centric environment. IDC believes that fabric-based network infrastructure called Ethernet fabric will dominate the discussion for the foreseeable future because it is an advance that enables organizations to build agile, business-ready datacenters and overcome obstacles such as infrastructure complexity.

Ethernet fabric is an emerging type of network designed to support the dynamic nature of cloud services in next-generation datacenters. The trend is toward a bigger, flatter network that preserves the low cost and simplicity of Ethernet, requires no reconfiguration or change in port configurations when resources move, and uses all paths available through the datacenter(s) for the most efficient transport of virtual machines (VMs) and their data. Therefore, enterprises seeking to optimize the investments they make in virtualization and cloud technologies should consider migrating to Ethernet fabric.

SITUATION OVERVIEW

Cloud computing is one of the prevailing IT trends as we head into the new decade. In fact, IDC predicts that cloud computing will move from a talking point to become one of the key transformation technologies in the marketplace. IDC believes that cloud computing will become the prevailing IT service delivery method as early as this year. By providing greater levels of automation, orchestration, provisioning, and deployment, cloud computing can help organizations become more nimble while they also reduce operating costs and improve resource utilization.

Cloud computing is emerging as a viable option for IT organizations to dynamically and automatically provision services without placing additional burden on their IT staffs. By enabling automated provisioning and deployment, the cloud model can help IT departments more rapidly scale their compute resources while they also retain the flexibility to adapt to changing business requirements.

IDC is forecasting ample growth in cloud computing over the next five years. A 2010 IDC survey of IT decision makers found that fully 44% are "considering private clouds," which reflects both user interest in cloud technology and the early stages of cloud adoption. Looking at the projected network buildout to support cloud deployments, IDC expects network infrastructure investments to more than double over the next three years, growing to over $1 billion each by 2013 for the public and private cloud segments (see Figure 1).

FIGURE 1

Worldwide Public and Private Cloud Network Forecast, 2009–2015

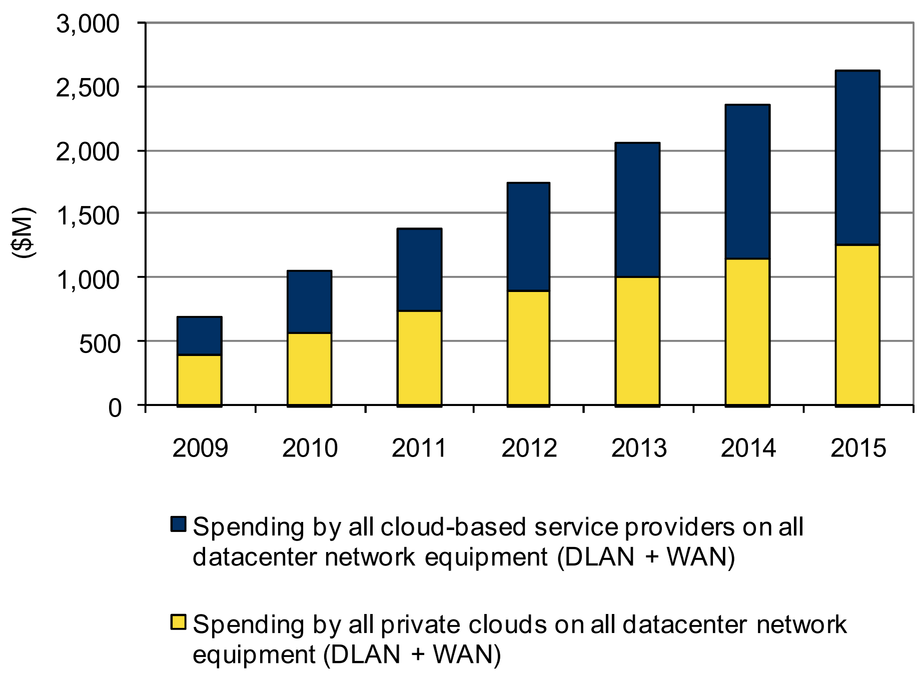

Notes:

Technologies included in the forecast are Layer 2–3 switches, Layer 4–7 switches, WAN application delivery, Fibre Channel switches, InfiniBand switches, and routers.

This forecast is a subset of the datacenter network forecast plus an estimate on the investment in routers to support datacenter communications (routers).

For more information on taxonomy questions, see *Worldwide Datacenter Network 2010–2015 Forecast and Analysis* (IDC #226224, December 2010).

Source: IDC, 2011

The allure of cloud computing lies in the dynamic service provisioning and resource allocation it enables and its "pay as you go" utility computing model. A key enabler of dynamic cloud services is virtualization. Virtualization technology delivers dramatic efficiencies in the datacenter by decoupling software resources from their physical hardware platforms using a middleware layer often called a hypervisor.

With hypervisors at work, network administrators can quickly create and move VMs — software instances of operating systems and application, Web, and storage servers — among compute nodes in the network as required. VMs use the available hardware processing, memory, and I/O resources of physical servers in a time-sharing manner, creating a flexible pool of available resources.

Virtualization offers valuable benefits for IT organizations, including the on-demand resource allocation described and the ability to sharply reduce the datacenter footprint. Server utilization has been reported to jump from as little as 15% when operating systems, applications, and computing hardware are bound together in a traditional fashion to as much as 85% using VMs, so it is understandable why virtualized environments have grown economically appealing. Datacenter and server consolidation lowers hardware equipment costs and related real estate, power, and physical management expenses. But the virtualization model also introduces new levels of complexity.

Current Datacenter Networks Inhibit Cloud Adoption

The increasing use of server virtualization on x86 systems has achieved cost and efficiency benefits, but realizing cloud computing's full spectrum of benefits will require more investment on the part of IT organizations. One area that requires significant attention is the network. Much of the virtualization activity to date has focused on the server and storage areas of the datacenter; to remain responsive to new IT demands and topology changes going forward, networks must evolve to keep pace.

Rigid Connectivity

For the network, a migration to virtualized IT means that all entities in the datacenter must be connected seamlessly. Currently, some networks are siloed to serve specific applications such as storage. More problematic, however, is that traditional datacenter networks face flexibility challenges caused by Ethernet's spanning tree protocol, the use of port oversubscription, and traditional tiered network architectures in which switches communicate up and down multiple switch "layers."

Ethernet's spanning tree protocol eliminates the undesirable traffic loops that would otherwise be created by redundant bridged Ethernet links that have been installed for failover and high availability. However, the protocol is slow to reconverge after a change or failure, degrading network performance. Aggregating links helps this situation, but in a virtualized environment, doing so limits the ability of a VM to move across physical servers unless cumbersome, manual port reconfigurations take place.

Oversubscription of ports becomes problematic because the client/server model flattens out in a virtualized environment. Traditional server-to-application ratios usually have been 1:1. In a virtualized environment, however, many applications might run on a single compute node. That means much higher bandwidth demands from multiple applications being made on far fewer physical ports.

Similarly, the "flat" model runs counter to the traditional three-tiered enterprise network designed with access, aggregation, and core switch layers. In the virtualized environment, the transmission of VMs and associated data is often peer to peer, server to server. So the typical configuration of traffic flowing up and down layers, across several "hops," introduces latency that can impact application performance. Such holdups run counter to the spirit of the cloud's dynamic resource provisioning benefits.

Limited Scale

The promise of cloud computing is to enable IT to proactively match IT resources (compute, storage, and network) to fast-paced, shifting business demands. For example, if an organization sees increasing business traffic at a particular time of day or by a specific business unit, it might be ideal to move the workload of a high-demand, congested compute node to another compute node that has more processing resources available. Unfortunately, traditional network architectures limit the scale and mobility of a workload. The movement of VMs adds further stress to the network. Performance, scale, and intelligence specific to each VM are required to meet the availability and resiliency requirements of the highly virtualized datacenter.

Virtualization and cloud computing needs are driving the industry to apply some of the characteristics of Layer 3 networking, such as link-state routing and equal cost multipathing, to Layer 2 in the form of Ethernet fabric and other virtual network models. This approach improves performance and avoids some of the slowdowns caused by routing table management.

Slow or Static Responses Undercut VM Mobility

Keeping the network in step with virtualized environments requires getting the dynamic, low-latency, lossless capabilities needed to quickly move VMs from one server to another, both within and between datacenters, as needed. If the network is slow or static, bogged down by spanning tree or multitiered datacenter Ethernet topologies, the network will impinge on resource provisioning dynamism.

Limited Ability for Policy Migration

The ease with which virtual servers can migrate within the datacenter for load balancing, maintenance, and energy efficiency is a cornerstone benefit of server virtualization. Currently, less than 10% of organizations have progressed to a fully virtualized, on-demand datacenter, where virtual workloads continuously and automatically migrate across the datacenter. The network is one of the barriers to moving along this IT maturity model.

The challenge arising from this process is how to keep connectivity and networking policy synchronized with physical location changes and to do so without manual intervention. IT managers need to retain network policy, service-level agreements (SLAs), and security policies in the face of VM migrations across different server resources and physical switch ports.

Need for Virtual Port Management

A network that is blind to virtual ports and that is tightly coupled with physical servers isn't sufficient for next-generation datacenters. In the past, traffic flows could be managed on a physical port basis. Now, they transcend physical ports as they pass from one virtual server to another and between virtual servers and virtual storage. The need for virtual port management is clear.

Additionally, VMs tend to be managed using server management platforms and tools. This creates a blind spot for network administration teams, making it difficult to effectively troubleshoot and manage VM network performance, particularly when VMs migrate across large network domains that cross several switches. This lack of visibility makes it difficult for network administrators to understand traffic patterns and create policies around VMs. Also, server virtualization software uses virtual switches that reside on the physical server to forward traffic between VMs on that server. This activity needs to be visible to all types of network switches.

Complexity

Complexity in the datacenter is the overarching inhibitor to cloud computing that arises from a network with a rigid architecture. Today, IT organizations have to invest in multiple layers of switches to guarantee bandwidth and performance. Continuing to add more equipment and capacity to classic Layer 2 Ethernet and Layer 3 routed networks can temporarily help the situation, but doing so also increases the complexity of the network, slowing down response times and increasing cost. It is important to reduce the number of components throughout the datacenter — specifically, the number of switch tiers — to reduce the number of times traffic flows up and down switch tiers so as to reduce latency.

Early Examples of Network Virtualization

To continue to keep datacenter services agile and dynamic will require full virtualization of datacenter networks. There have been several examples of network virtualization deployment over the years, such as segmenting Ethernet LANs into virtual LANs (VLANs), creating encrypted virtual private networks (VPNs) across a shared WAN service, and even running virtual routing and forwarding (VRF) instances within a single routing appliance.

These flavors of network virtualization have been invaluable for partitioning data among user groups in the LAN, for ensuring security and privacy in the WAN, and for reducing router appliance requirements at the enterprise WAN edge. But they do not address the scalability and dynamic requirements of the cloud. To truly enable the easy adoption and management of a virtualized datacenter and help address the bottlenecks described earlier requires an Ethernet fabric–type approach.

Ethernet Fabric Enables the Next-Generation Datacenter

IDC believes that customer demand to migrate to agile cloud computing will lead to new intelligent networks that support the notion of fabric computing. Cloud-optimized networks will enable resources to move freely through the network and at the same time will reduce costs and complexity. An on-demand datacenter is critical to helping a company increase business agility and respond to revenue-generating activities much more quickly. Ethernet fabric networks are one example of networks emerging to meet this goal.

The fundamental idea behind Ethernet fabric is to create a more efficient, cost-effective way to move data through a network. Ethernet fabric networks are switch networks designed for greater flexibility and simplicity than today's Layer 3 routed networks and for greater scalability than classic hierarchical switched Ethernet networks based on spanning tree protocol. Ethernet fabric involves what could be loosely described as "Layer 2 routing." The ability to build large, fast datacenter networks that can be "routed" based on location-independent Layer 2 MAC addresses has become possible now that memory and processors have grown far more powerful and less expensive than they were when Layer 3 routers were first designed.

Fully virtualized datacenters built with an Ethernet fabric offer the following benefits:

- **Inherent scalability.** The amount of data managed by the enterprise is projected to grow exponentially over the next few years, putting even greater demands on already overloaded networks. IT managers will need to implement easily scalable networking infrastructures to be prepared to manage the accelerated growth of data and network traffic. Ethernet fabric solutions combine the best attributes of traditional Layer 2 switching (speed) with those of Layer 3 routing (intelligence and multipath data forwarding) to determine best paths for transmissions and to preserve policies as VMs move around the datacenter.

- **Sustainable performance.** In addition to preserving VM policies, the speed and intelligence attributes described enable full link utilization with redundancy, deterministic multipathing, and subsecond link reconvergence. As a result, VM performance is optimized while still enabling high-availability networking across redundant links.

- **Seamless VM mobility.** It is currently fairly simple to move VMs within a blade chassis, but it becomes more difficult as the VMs' sphere of mobility expands across chassis, across racks, and across datacenters. The network, then, needs to be virtualization aware such that port policies follow VMs as they move to different compute nodes.

- **Logical management.** The continuing desire to create and enable efficient IT operations is bringing network management to the forefront. Operational simplicity is being achieved with the introduction of virtual chassis architectures that provide the ability to access multiple switches while logically managing those switches as one device. Ethernet fabric enables physical switches to be logically collapsed into what appears to network administrators as a single "big switch" — or a pool of switching resources.

- **Increased VM visibility.** VMs are also becoming a shared management issue across IT domains. Ethernet fabric provides network operations teams with greater visibility into how VMs are behaving on the network. The sheer number of entities to manage in a virtualized datacenter, where VMs can be provisioned in minutes, grows swiftly and can get quickly out of hand.

- **Automation of VM mobility.** The flexibility inherent in Ethernet fabric simplifies the ability to provision and deprovision VMs either on demand or in accordance with business rules and SLAs.

The ability to simplify and consolidate both physical cabling and the number of devices in the network will drive customer adoption of Ethernet fabric. The overarching mandate is for enterprises to make sure that all components of their virtualized datacenters and cloud services — including the network — are in sync so that no single component lags and inhibits migration. The single greatest change network managers can make in their datacenters today that will enable them to achieve their goals is to migrate to an Ethernet fabric for reduced infrastructure complexity, increased IT flexibility, and improved workload mobility.

DRIVING AN ETHERNET FABRIC FOR THE CLOUD WITH BROCADE VDX SWITCHES

Brocade recognized the new requirements of virtualized cloud networks early on. In response, it has embedded its Virtual Cluster Switching (VCS) technology into Brocade VDX 10GbE switches to deliver an Ethernet fabric.

Brocade's goal for a VCS-based Ethernet fabric is, first and foremost, to provide application availability and resiliency. Additionally, Brocade aims to provide a fabric that reduces infrastructure complexity while providing agility to respond to changing business requirements.

Brocade VCS/VDX offers the following attributes:

- **Deterministic multipathing and low latency.** By leveraging multiple paths through the network and continuously determining the most efficient route, the fabric enables high performance while maintaining high availability. Port-to-port switching is as low as 600 nanoseconds, so latency is extremely low.

- **Distributed intelligence.** VMs can move seamlessly and transparently to any compute node in the datacenter or another datacenter. Distributed intelligence also enables the automatic migration of port profiles to ensure that network policies automatically follow VMs when they move.

- **A logical chassis.** The fabric appears to network administrators as one large switch and delivers increased traffic visibility and better control of the network while reducing network management overhead. While all switches in an Ethernet fabric are managed as if they were a single logical chassis, they appear no different from any other single Layer 2 switch to the rest of the fabric. Each physical switch in the fabric is managed as if it were a port module in a chassis. This enables fabric scalability without manual configuration; when an additional physical switch connects to the Ethernet fabric, all applicable fabric settings, including interswitch link (ISL) configurations, are automatically created with no manual intervention necessary.

- **Dynamic services.** Network administrators will soon be able to insert services such as encryption (security), load balancing, and fabric extensions and reconfigure the network dynamically without physical rewiring or experiencing downtime.

- **Support for datacenter automation.** The distributed intelligence in Ethernet fabric environments aims to improve VM management by automatically migrating port profiles along with the VM when it moves through the datacenter. This capability is hypervisor agnostic, supporting the use of more than one hypervisor technology in the datacenter. As mentioned, classic Ethernet network policies are static, associated with a particular physical port, so when a VM moves, its traffic shows up on a different switch port, leaving behind the static policy on the initial physical port. The result is that the policy is not applied to the VM on the new switch port. A distributed control plane database in Ethernet fabric, however, makes policies accessible to every port in the fabric. When a switch port sees traffic from an application for the first time — the MAC address of its VM — it checks the distributed policy database for a policy matching that MAC. If there is one, the fabric applies it.

- **Seamless interoperation with classic Ethernet networks.** The VDX's Ethernet fabric seamlessly integrates into traditional Ethernet switched networks. VDXs can be used as traditional Layer 2 10GbE switches for now, and enterprises can simply "turn on" the Ethernet fabric capability when they are ready to virtualize their datacenters.

With VCS, common configuration parameters and device information are automatically distributed to each member switch in the fabric. When a server connects to the fabric for the first time, all switches in the fabric automatically learn about that server. Also, when two VDX switches are connected, the fabric is automatically created, and the switches discover the common fabric parameters. This fabric configuration is shared among all switches in the fabric, rendering it masterless, so that no single switch stores configuration information or controls fabric operations.

CHALLENGES AND OPPORTUNITIES

Because enterprises will want to optimize their investments in virtualization and cloud technologies, they will ultimately seek a network that adequately supports the new environment. As increasingly sophisticated virtualization use cases emerge with VM mobility as their foundation, enterprises will look to enhance their Ethernet and storage networks to meet the new scalability and automation requirements.

The biggest challenge will be for Brocade and other Ethernet fabric providers to articulate the benefits of using flatter networks with distributed intelligence rather than traditional Layer 2 switched or Layer 3 routed networks because of their static nature. In addition, it will be important for Brocade to foster a complete ecosystem to work with a full range of higher-level datacenter automation and provisioning systems.

It will also be important for vendors such as Brocade to educate enterprise buyers about Ethernet fabric's ability to seamlessly integrate into traditional Ethernet switched networks. Prospective buyers should know that Ethernet fabric switches can be used as traditional Layer 2 10GbE switches and, as such, do not require enterprises to fully rip and replace their existing switching infrastructures. Still, coexistence will be temporary, as eventually it will be desirable for all virtualized servers and storage to connect to the Ethernet fabric for all the reasons discussed in this paper. At that point, customers can simply activate the Ethernet fabric capability in their switches.

CONCLUSION

To reduce the cost and complexity of their IT infrastructures while creating an on-demand datacenter that can facilitate time-to-market demands and revenue growth, enterprises will continue to invest in the datacenter. Top of mind to reach this goal is to consolidate datacenters, virtualize resources, and embrace cloud computing. IDC believes that as part of these processes, enterprises wanting to reap the full benefits of their consolidation, virtualization, and cloud investments must also soon begin to virtualize their datacenter networks.

A fully virtualized network is necessary because the VMs that help enable the attractive on-demand provisioning capabilities of cloud computing are no longer tied to physical switch ports — so neither are network management and security policies. With VMs hopping among physical computing resources as those resources become available, enterprises need a way to keep policies associated with those VMs other than by static physical port numbers.

IDC believes that this requirement will drive enterprise investments in network virtualization using a newer, cloud-optimized model called Ethernet fabric. The cloud-optimized Ethernet fabric network will have the following characteristics:

- **Flat and fast.** The networks are intelligent enough that switches in the fabric know about one another and can find the shortest transmission path without creating inefficient loops.

- **Quickly scalable.** The networks will scale quickly and easily because all ISLs are active and available. Traffic automatically travels along the shortest path, reducing the latency involved in communicating among devices up, down, and across the hierarchy of switch layers.
- **Simpler to manage and secure.** The fabric is managed as a single logical entity, allowing policies and profiles to be applied globally and to automatically follow VMs across physical servers and storage infrastructure.

These attributes will enable customers to build an agile, business-ready, and cloud-ready datacenter while overcoming infrastructure complexity obstacles.

Copyright Notice